

The Automated Planet Finder at Lick Observatory

Matt. V. Radovan^a, Kyle Lanclos^a, Bradford P. Holden^a, Robert I. Kibrick^a, S. L. Allen^a, William T. S. Deich^a, Eugenio Rivera^a, Jennifer Burt^a, Benjamin Fulton^b, Paul Butler^c and Steven S. Vogt^a

^aUniversity of California Observatories, 1156 High St., Santa Cruz, CA, USA;

^bInstitute for Astronomy, University of Hawaii, Honolulu HI, USA;

^cCarnegie Institution of Washington, Department of Terrestrial Magnetism, 5241 Broad Branch Road NW Washington DC 20015-1305 USA

ABSTRACT

By July 2014, the Automated Planet Finder (APF) at Lick Observatory on Mount Hamilton will have completed its first year of operation. This facility combines a modern 2.4m computer-controlled telescope with a flexible development environment that enables efficient use of the Levy Spectrometer for high cadence observations. The Levy provides both sub-meter per second radial velocity precision and high efficiency, with a peak total system throughput of 24%. The modern telescope combined with efficient spectrometer routinely yields over 100 observations of 40 stars in a single night, each of which has velocity errors of 0.7 to 1.4 meters per second, all with typical seeing of < 1 arc second full-width-half-maximum (FWHM). The whole observing process is automated using a common application programming interface (API) for inter-process communication which allows scripting to be done in a variety of languages (Python, Tcl, bash, csh, etc.) The flexibility and ease-of-use of the common API allowed the science teams to be directly involved in the automation of the observing process, ensuring that the facility met their requirements. Since November 2013, the APF has been routinely conducting autonomous observations without human intervention.

Keywords: Automated Telescopes, Early Operations

1. INTRODUCTION

1.1 Science case for the APF

To date, the vast majority of planets discovered around other stars, or extra-solar planets, have been detected via indirect observations. This requires measuring the effect of the planet on the observable properties of the host star. Because planets are much smaller and less massive than the stars they orbit, detection requires searching for small changes, fractions of one in ten thousand, in the host star's flux or radial velocity. In order to isolate the signature of a planet from a background of other random processes, frequent observations of the same star are required with a stable and reliable instrument.

Doppler spectroscopy, or the radial velocity method, is one of the earliest methods to successfully detect an extra-solar planet. By measuring the Doppler shift in a star's spectrum resulting from a planet's pull, the orbit of the planet may be calculated; with refinement over the last several decades, changes of only a few meters per second can routinely be detected with most high resolution spectrometers. Pushing the sensitivity below a meter per second opens the possibility of detecting habitable planets around sun-like stars.

The combination of the requirements of stability and a high cadence of observations naturally suggests that a dedicated, automated facility for radial velocity searches would be at the forefront of extra-solar planet discovery. A dedicated facility would ensure stability, as there would be no need for instrument and telescope changes that might occur at a more general use facility. A dedicated facility also allows scientists to meet strict cadence requirements without impacting the science programs of other astronomers. As the constant observing would be a severe drain on participating astronomers, automating the routine nighttime operations is a practical requirement for such a facility. These conclusions became the basis of the Automated Planet Finder¹ at Lick Observatory, located about 32 kilometers east of San Jose, California, and shown below in figure 1.

Further author information: (Send correspondence to B.P.H.)
B.P.H.: E-mail: holden@ucolick.org



Figure 1. APF telescope at twilight, with the lights of San Jose in the distance

1.2 Major components of the APF

1.2.1 The APF Dome

The dome for the APF facility is an IceStorm-2 enclosure, designed and constructed by Electro Optical Systems (EOS).² The dome is slaved to the azimuth axis of the telescope, and will track any/all motion in azimuth so long as the dome is powered. The dome slit has two independently-movable shutter panels, which can be ‘parked’ in an up-and-over mode that nearly parks them out of view of the telescope, or in arbitrary positions to provide wind shielding for the telescope.

1.2.2 The APF Telescope

Designed and constructed by EOS Technologies (EOST), the telescope has a 2.4m f/1.5 primary diameter mirror on a Alt-Az mount. To maximize optical collecting area, the secondary mirror and its support structure were kept small, such that only 2% of the central area was obstructed. One consequence of a small secondary is a reduced field of view, which for the APF is roughly one arc minute. This implies a requirement for the blind pointing of the telescope, which has a root-mean-square (RMS) pointing error of less than three arc seconds and a worst case pointing error of about 12 arc seconds. The telescope nominally slews at 2-3 degrees per second, a relatively high speed which reduces the loss of observing time while moving to a new target.

1.2.3 The Levy Spectrometer

Designed and constructed by University of California Observatories (UCO), the Levy Spectrometer is engineered for high resolution spectroscopy. The Levy is mounted at the left Nasmyth port of the APF telescope, attached directly to the telescope yoke; the spectrograph thus never changes orientation with respect to gravity, and the input beam remains stable as the telescope tracks in azimuth. The wavelength calibration for the Levy is provided by an iodine cell³ inserted into the focused beam after it enters the spectrometer. The total system throughput is 15% when the iodine cell is in place; though the APF has only 1/12 of the light gathering surface as the 10-meter Keck 1 telescope, the high efficiency of the Levy spectrometer means that exposures on like targets are only 6 times longer than similar exposures taken with the HIRES⁴ spectrometer on Keck 1.

2. ROADBLOCKS TO ENABLING APF SCIENCE

Site acceptance of the APF telescope occurred in November 2010; the Levy spectrometer was delivered in January 2011. Final commissioning began in June 2011, after a delay resulting from a failure of the space frame structure that occurred during shipping of the instrument.⁵ Three major areas of work were addressed during commissioning: the stability of the Levy spectrometer,¹ the tracking performance of the APF telescope,⁶ and the development of both procedures and infrastructure to remove the human element from routine nighttime operations. The first two topics are summarized here, as they are addressed more thoroughly in the cited publications; the third topic represents the bulk of this paper.

2.1 Stability of the Levy spectrometer

The stability of the spectrometer is measured by its ability to maintain a stable point spread function and precisely and consistently place on the science detector specific emission or absorption lines from a calibration source. In the ideal case, such lines will always land in precisely the same physical location on the detector. This stability not only reduces the number of variables to consider when comparing data between nights, but it also ensures that the calibration spectra taken at the beginning and/or end of the night are appropriate for the exposures of science targets. The use of an iodine cell to provide wavelength calibration largely eliminates the first concern, but it cannot eliminate the second, as not all calibrations are wavelength-centric; for example, the focus of the spectrometer has a temperature dependence, and variation in the focus will not be corrected by the use of an iodine cell. To test the stability of the Levy spectrometer, continuous sequences of exposures of a thorium-argon emission spectrum were performed under a variety of conditions.

Because the spectrometer is not vacuum-sealed, the spectrum is subject to shifts in the image plane as a result of changes in air pressure. Careful analysis of the stored meteorological telemetry,⁷ when combined with the positions of the thorium-argon emission lines, showed that larger shifts in line position were correlated with changes in pressure rather than temperature. At the same time, significant shifts in line position resulted from opening and closing the fiberglass spectrometer enclosure. Further investigation found several mechanical interference points between cabling and hoses on the instrument's optical bench and the fiberglass enclosure; removing the interference nearly eliminated the correlation between changes in pressure and shifts in line position. The remaining correlation is tolerated, with the assumption that poor seeing will overwhelm any shifts resulting from a large overnight change in air pressure.

Changes in line position as a result of temperature changes are, to first order, compensated for by the support structure for the major optical elements in the spectrometer. The support frame has thermally sensitive components that will passively piston and rotate optics to counter shifts in the dispersed spectrum resulting from temperature changes and the limitations of near zero coefficient of thermal expansion materials like Zerodur and Invar.^{1,8} To further mitigate the effect of temperature changes, the spectrometer is enclosed in an insulating fiberglass shell; the interior of the spectrometer is then maintained at a controlled temperature year-round. One important reason for doing so is to minimize the magnitude of residual errors from the passive compensation system and to reduce the effects of significant uncompensated components such as the steel lead screw on the CCD focus stage. Identification and mitigation of thermal shorts to the spectrometer interior reduced the load on the instrument chiller and on the heating pads. Doing so provided the necessary thermal capacity to maintain the target temperature in the depths of winter and the heights of summer, with average nighttime outside air temperatures shifting more than 20 degrees Celsius; precedent also exists for temperature swings on the order of 10 degrees Celsius over the course of a single night.

A more thorough summary of the stability of the Levy spectrometer is presented in Vogt's paper¹ and a detailed discussion of the mechanical design by Radovan.⁸ The last steps not directly mentioned in that paper involved severing the last few correlations between the spectrograph interior and the dome interior. The most significant change was to relocate the heating pads for the interior of the spectrometer; though the pads were covered with optical masking material, they were sufficiently within line of sight of the camera and/or mechanical support for the dewar to introduce a shift in the detector plane. At the same time, additional circulation fans were installed inside the enclosure to further improve the mixing of ambient air around all spectrometer components. After making these improvements, the absolute focus of the spectrometer now remains consistent within fractions of a micron on a night-to-night basis regardless of exterior conditions.

2.2 Telescope tracking performance

When the APF telescope passed site acceptance in November 2010, its tracking performance met or exceeded its required performance specifications; at that time, the Levy spectrometer had not been installed, though a system of dummy weights was present to emulate the weight distribution of the as-designed spectrometer. Over the next year, the dummy weights were replaced with the spectrometer, which necessitated removing several covers on the telescope yoke and the installation of significant cabling. With the APF in its final configuration, it was clear that the telescope control system was inducing significant oscillations in the motion of the telescope, effectively rendering the telescope unable to maintain the position of a star inside the boundaries of an aperture at the focal plane. Beyond poor performance, the telescope also experienced uncontrolled oscillations, which prompted legitimate concerns for the safety of the facility itself. Later efforts determined that the telescope is very sensitive to increases in its rotational inertia about the azimuth axis; in hindsight, it seems likely that the changes in the mechanical configuration were sufficient to push the overall control system system beyond the region of stability offered by its as-delivered configuration.

The first problem (poor tracking) reduced the efficiency of observing time and unevenly illuminated the aperture at the focal plane, introducing an error in the attainable radial velocity precision. The second problem (uncontrolled oscillations) prevented the use of the facility in an unattended mode, as a human was required to intervene to protect the health of the facility. Thus, the poor telescope control performance effectively shut down the facility.

While the telescope was still under warranty, we went through the vendor to resolve these issues; when the warranty expired without resolution, UCO began an internal initiative in September 2012 to resolve the outstanding problems. That initiative succeeded in June 2013, at which point the facility immediately began taking world-class data. A detailed account of the problems encountered and solutions applied is provided in Lanclos' paper,⁶ also in this conference.

2.3 Automating routine operations

The APF telescope, as delivered, had most if not all of the necessary infrastructure to operate the facility with a human at the controls. The assumption of direct human involvement permeates many of the design decisions. A pair of quick examples: at a high level, we were encouraged by EOST to take data for, calculate, and install a new pointing model on a monthly basis, with all observations performed by hand; at a low level, the default response of dropping power to the entire telescope as a result of any failsafe violation anywhere in the facility, making error recovery unnecessarily difficult, and remote recovery impossible. The challenge, which began well before the telescope received site acceptance, was to develop sufficient infrastructure around the as-delivered control systems to automate routine observing tasks.

As the infrastructure took shape, and as the observing teams began to empirically determine what was required to successfully work with the as-delivered facility, additional development went towards streamlining and automating routine tasks. As tasks were automated, procedural questions would arise, especially with respect to ensuring personnel and facility safety. One of the final developments was each science team's creation of its own master-of-observing software, implementing their specific preferences for how a full night's observations should be conducted. Barring the occasional equipment failure, the APF has been operating in a near-autonomous mode continuously since November 2013.

The remainder of this paper addresses in detail some of the specific problems encountered during this period of final commissioning and early operations.

3. AUTOMATING HUMAN-ORIENTED CONTROL SYSTEMS

The control software delivered by EOS and EOST (hereafter EOS(T)) included a set of control daemons, graphical user interfaces, and a C++ library to allow access to the inter-process communications used by these software components. Successfully using the facility involved a lot of manual data entry; successfully recovering from errors universally involved restarting major software components, and/or rebooting control hosts. This approach may be sensible, perhaps even cost-effective, if the intent is to always have a technician at the controls. Attended observing at the APF was seen strictly as a transitional state of affairs, as opposed to a long-term design goal;

the software stack from EOS appears to hold to the reverse, that attended observing was the design intent, and any infrastructure for automation came about only as a means to that end, rather than being an end in and of itself.

The general practice at UCO when dealing with discrete hardware systems is to directly expose all of the telemetry provided by a given hardware controller. If exposing everything is not practical, a useful subset is instead provided, but with the same key caveat that the exposure at the lowest level is direct and unfiltered. For example, a dome shutter may expose its absolute position, current velocity, motor torque, and a host of other parameters. From this fine-grained access, higher level functionality can be constructed: when trying to move two dome shutters to individual positions, the appropriate command sequence will vary depending on the current state of the system, and the assessment of success will likewise vary depending on the actual command issued. Proper handling of these system-level objectives can thus be fully described in a layer that builds upon the telemetry exposed at the lowest level. Successfully adapting the EOS(T) control software required both of these approaches.

3.1 Establishing low-level access

EOS provided a Linux-compatible API for the dome and telescope control system software as part of the contract for the overall facility. In order to make the best use of the standard infrastructure at Lick Observatory, immediate work began to wrap the vendor-provided API inside the Keck Task Library (KTL)^{9,10} API, with the intent of directly exposing every available parameter via KTL keyword/value pairs.

At the time, UCO was the only customer for this API, certainly the only customer for Linux-based access; this manifested itself in several ways. The vendor API changed several times. Old function calls for setup and access to the telemetry were replaced by new strategies, signatures of function prototypes changed incompatibly, interface specifications did not describe the meaning of parameters, underlying software did not provide some documented parameters, and the telescope subsystems did not implement new features of the API. EOS was unwilling to include the source code as an immediate contract deliverable; instead, all source code was placed into an escrow account, with a planned release of escrow in 2018, eight years after initial site acceptance. The absence of this resource proved to be a significant stumbling block, with unanswered questions and inconsistent behavior persisting despite years of local effort empirically determining how to handle access to the EOS client API.

It is perhaps because the low-level access is so critical to the success of all the software built on top of it that these difficulties were so visible. Despite the lingering errors and inconsistencies, this effort to insulate the foreign API inside what is for Lick Observatory a universally-understood API remains a key component enabling the success of the overall facility.

3.2 Establishing high-level subroutines

Having succeeded in establishing low-level access, effort turned towards providing high-level access to complete systems: slewing the telescope, manipulating the shutters, opening and closing vent doors, etc. The commands provided by the vendor employ a fire-and-forget strategy where the only direct feedback is to acknowledge whether the command was received with valid syntax. Ascertaining whether the underlying software and hardware successfully performed the command generally requires monitoring numerous other parameters whose relationships are not part of the documentation.

In order to retain the unfiltered nature of the low-level access, no effort is made to take into account any idiosyncracies in the behavior of the underlying control system. To use an example, a boolean value is provided by the dome control software to indicate whether the dome shutter is closed; the precise logical formula that determines the value of this parameter is not known, so instead we rely on empirical observations of the parameter as compared to other, related parameters in the system. Most of the time, when the dome is open, the boolean parameter for dome closure indicates 'False'. Most of the time, but not all of the time; properly handling these inconsistencies became a major role of the software layered on top of the low-level access, so that each individual tool that needs to manipulate a given system does not have to reimplement complex rules to take into account a wide variety of potential failure modes.

That these high-level subroutines can in turn expose their functionality as KTL keyword/value pairs is a deliberate and important design decision. From the perspective of a client application, the act of parking the telescope has the same complexity and failure modes as setting a temperature setpoint; a client does not need to be aware that the park operation has to query dozens of parameters as part of its command sequence and sanity checks, but at the same time, the application does not need to limit itself to high-level parameters: the full breadth of the parameter space is available to every application at all times.

The combination of direct access to all telemetry and the aggregation of both telemetry and error handling into modular, system-level interfaces enabled the pursuit of the next layer of software (see section 5), sequencing the actions of both discrete and integrated systems to perform observational activities.

4. ENSURING PERSONNEL AND FACILITY SAFETY

We used our experience with remote observing at Lick Observatory’s Nickel 1-meter telescope¹¹ to shape our approach towards first enabling remote observing with the APF, and then on to enabling autonomous nighttime operations. This included some re-use of existing software, but more important was the re-use of lessons learned.

One key experience with enabling remote observing at the Nickel was that most observers quickly abdicated their role in monitoring conditions once software was provided to assist in that process. In response, that same software evolved to take full responsibility for monitoring conditions and closing the dome as necessary. A smaller number of observers tried to disregard the published observing limits and keep the dome open in inappropriate conditions. This experience led us to adjust the oversight software to adopt the same strict, conservative approach whether the facility is operating with local or remote observers. This directly informed the approach taken for the APF, and led us to be stricter still when operating autonomously. Taking Doppler spectroscopy data lends itself well to a restrictive approach, in that one observation represents a small piece of the full data set for that star, which may comprise upwards of one hundred individual observations; rescheduling the observation for a later time usually does not impact the overall quality of the science. This is very different from other observing modes, where a single observation may be the full data set for a given target.

4.1 Permissions for discrete operations

One software package has responsibility to monitor all aspects of the APF that affect personnel and facility safety. All conditions are reduced to three simple permission states: whether opening the dome is allowed, whether moving the telescope is allowed, and whether operating the Levy spectrometer is allowed. Other applications need only check these three constraints before attempting a potentially restricted operation.

Striking a balance between personnel and facility safety and the desire to maximise scientific output is key. The permissions must be appropriate to the actual use of the facility, and not interfere needlessly with nighttime operations; if too many false positives occur, our observers would surely begin to subvert the permission scheme as a practical means towards achieving their observational goals. While recognizing the need for balance, the permissions must be structured such that the default action is to prevent autonomous control systems from creating hazardous conditions for personnel entering the dome, and the facility itself must be protected from hazardous environmental conditions.

4.2 Honoring permissions

As described in section 3, all operations can be expressed at their lowest level as one or more KTL keyword/value pairs. Many of these keywords support both read and write operations, and in the case of mechanical systems, writing a value to a keyword implies a command to physically move the mechanism. It is neither realistic nor good practice to expect that every application will know and apply the correct permission for each and every one of its KTL keyword modification requests. Nor does the KTL architecture provide any assistance in restricting such requests, such as mandatory acquisition of rights or similar features. To work around this deficiency, we took the simple approach of replacing the standard keyword-modification interface with a “safe” version. The safe version performs a table lookup of constraints that apply to the keyword, queries the current values of any such constraints, and only if all of the permission checks pass is the original modification request issued.

This *preventative* approach for honoring constraints is one key usage of the established permission scheme, but it is only half of the equation; while all individual subcomponents are expected to honor the overall permission structure, we cannot assume that this will be the case. Thus, the preventative initial approach is combined with two flavors of reactive sentinels running in the background. The first sentinel is internal to the permission system itself: if it sees that a severe condition has arisen, such as rain falling on an open dome, the permission system will immediately and repeatedly issue commands to close the dome, ignoring the possibility that other software may be attempting actions to the contrary. The second reactive sentinel integrates with the overall task structure (see section 5); this second sentinel will immediately shut down any task whose critical permission has been revoked.

The combination of both a preventative and a reactive approach has been very successful in ensuring both personnel and facility safety at the APF.

4.3 Monitoring arbitrary conditions

Where the permission rules are kept deliberately simple and sequestered from the rest of the software infrastructure, there is also a need to create conditions of arbitrary complexity, providing alarms for undesirable states, automatically correcting anomalous conditions, or creating derived data products. One simple example of the latter is the average strut temperature for the space frame inside the Levy spectrometer; a more complex example is a thermal control loop that attempts to maintain that same average strut temperature at a constant value. Both of these examples, and many, many others in between, are implemented using an in-house software monitoring tool.¹²

The basic premise of this monitoring system, as with any generalized monitoring tool, is that any undesirable condition should be watched for, and appropriate actions taken in the event that such a condition occurs. The system was initially configured to track system performance with respect to documented requirements, such as telescope tracking accuracy, or iodine cell temperature; all hardware failure modes that could be identified in advance, such as weather station availability, or failures in the battery backup system; and unexpected observing configurations, such as incorrect settings for the guide camera. The set of monitored conditions grew steadily over time as additional problems and failure modes were discovered empirically, from an initial set of about 100 conditions, to about 300 conditions as of this publication.

For conditions that can be corrected automatically, this represents a direct and complete resource savings, in that no direct intervention or acknowledgement was necessary to remedy a problem; one such action ensures the operational safety of the secondary mirror, disabling motion of the mirror actuators if the tip/tilt of the secondary exceeds a rational limit. For conditions that require human intervention, the time must still be spent to correct the problem, but it was not necessary to wait for the failure to become critical or otherwise have an observable negative impact on the performance of the facility before it was noticed; an example may include notification that persistent software has unexpectedly exited, and needs to be restarted. As the APF extends its operational lifetime, the cost benefits of implementing a generalized monitoring system can only increase.

5. MUTUALLY-AWARE DIVERSE, AUTONOMOUS TASKS

One key challenge for the next layer of automation was not technical at all, but rather budgetary, in that major aspects of software development could not be continued by the software engineering group, instead becoming the responsibility of at least four different parties with a vested interest in the success of the facility. In order to successfully integrate their efforts, and to improve our ability to monitor ongoing activities at the APF, it was necessary to establish a framework that was simple enough to present a low barrier for all authors, yet at the same time provide enough functionality that there was a material benefit for the broader objective of integrating their efforts.

5.1 A common framework for diverse tasks

From an external perspective, the key requirement is that these diverse tasks must be represented by the same fundamental telemetry. By doing so, this allows all other software concerned with the scheduling and querying of tasks to be presented with a uniform interface, regardless of the nature of the task being performed. Some

tasks are transient, some run at all times; many tasks concern themselves only with a specific function, but other tasks will invoke multiple sub-tasks.

As mentioned in section 4, permission constraints are handled in both a preventative and reactive sense. This task framework operates in both ways: it prevents multiple instances of a task from running simultaneously, and provides the necessary hooks to allow sentinel applications to explicitly shut down tasks if their key permissions have been revoked. These two fundamental use cases, eliminating duplication and forced shutdowns, drove the structure of the basic task representation in KTL keywords. In broad strokes, the task framework tracks whether and where a task is running, what permissions are required by the task, and how the task would prefer to be signaled in order to initiate a shut down. In addition, the task framework allows arbitrary KTL keywords to be added for each task, which allows the task implementations to make use of custom, task-specific KTL keywords without the need to reimplement the task to be fully KTL-aware. By externalizing the KTL implementation of the task keywords, this reduced the development load on the task authors.

5.2 Simplifying the adoption of a common framework

The original intent of the task framework was to implement much-needed functionality without increasing the development burden on the task authors; in the ideal case, the development burden could actually decrease, in that the task framework could transparently handle some amount of complexity without directly exposing it to the task author.

Task implementations were written in a diverse set of scripting languages, some without native language support for KTL. Rather than restrict the language choices, we instead implemented a simple command-line interface that task authors could use to interact with the overall task framework. Beyond the basic functions of handshaking with other running task instances and manipulating KTL keywords, the command line interface also transparently implements key aspects of the signal handling mechanism. Specifically, tasks can be told to proceed with normal execution, to pause execution, or to immediately abort execution. Any time that a task would perform an operation that waits for another condition to occur, such as suspending itself for a few minutes, or waiting for a KTL keyword to transition to a specific value, the task implementation will use the task command line interface to perform this blocking operation; the command line interface will simultaneously track the KTL keywords that indicate whether it needs to pause or abort, and take action automatically. Because the task command line interface allows the execution of arbitrary commands, this automatic functionality can be leveraged for any and all extended operations performed by the script.

6. IMPLEMENTING AUTONOMOUS OPERATIONS

The coarse outline for the activities in a 24-hour period are the same for many observational facilities, and APF is no exception: calibrations are performed in the late afternoon; the facility is prepared for observing; observations occur throughout the night; the facility is safely stowed; calibrations are performed at the conclusion of a night's observing. With safety concerns (section 4) and idiosyncracies largely handled at lower levels (section 3), the implementations for these steps could focus on the subject at hand.

6.1 Scripted Observing Tool

Taking a single observation touches on nearly all aspects of the facility. As such, the automation of observing is substantially more complex than the infrastructure developed for individual systems, such as the moving the dome shutters or slewing the telescope. The implementation must also be efficient, as any unnecessary delays manifest directly as lost observing time.

The initial implementation of our scripted observing tool was successfully used in May 2013 by one science team, and by a second science team in June 2013. What started as a simple, quick-and-dirty prototype to provide a succinct interface for scripted observing evolved organically over the next 6 months in response to feedback from the observing teams.

The scripted observing tool limits its scope to performing the sequence of steps necessary to take an observation, or sequence of observations; it does not attempt to do any dynamic scheduling or to otherwise optimize the selection of targets or the order in which they are observed. It also does not monitor the weather or many other

failure modes, instead relying on other software components (see section 4) to provide appropriate notification should conditions exist that would prevent observing.

The scripted observing tool will process the entries from the target list one line at a time until either: all of the entries in the star list have been processed, the script determines the elevation angle of the sun is > -8.9 degrees, the script encounters an unrecoverable hardware or software fault, or if the tool receives an explicit signal to stop execution. Lists of any length are allowed, including a list with only a single target. This latter case enables the use of dynamic observation schedulers, which select the best available target for the present conditions.

Prior to acquiring each target, the scripted observing tool performs a trajectory analysis for that observation to ensure the target's accessibility, to select the correct portion of the azimuth cable wrap and to determine the position of the dome shutters that will ensure the observation will not be vignetted while also minimizing the exposure of the telescope to the wind. Then, the tool slews the dome shutters and the telescope in both axes in a manner that protects the primary mirror while minimizing the overall slew time. Once the target star is in the field of view of the guide camera, the scripted observing tool dynamically adjusts the guide camera's exposure parameters to compensate for variations in target brightness and sky transparency. It then centers the target on the spectrometer slit, enables guiding, and initiates the requested sequence of science exposures on the specified target.

Over the course of initial operations, two additions to the scripted observing tool were made. First, an option was added to allow observing in crowded fields or cases where there may be difficulties in target acquisition. The worst-case pointing for the APF is still good enough to enable direct acquisition of most target objects by doing a blind slew to the target's coordinates and then locking onto the brightest object in the guide camera's field of view, which is roughly one arc minute. However, for targets with nearby neighbors that are either brighter or of comparable brightness, an additional step is needed to reduce the local telescope pointing error to ensure that the target object is correctly identified and acquired. To handle this case, a nearby reference object is selected from a catalog of bright stars in uncrowded fields. The telescope is then slewed to the coordinates of that reference object, and the autoguider is commanded to center up on the reference object. The telescope is then offset from that reference object to the actual target. This minimizes most pointing errors to less than two arc seconds, which allows the correct science target to be selected from the guider field.

The second addition allows the telescope to be focused prior to beginning a science exposure on a target. By default, if the target is bright enough that the guider is set to a frame rate of 1 Hz or faster, the software will check the current telescope focus by pistoning the telescope secondary by a fixed amount in both directions. If that test finds that the focus has changed, then four additional offsets centered on the previous focus position will be used to calculate the new best focus. The user, when assembling the list of targets, has the option to force the full seven-step focus loop before the observation.

6.2 Nightly scheduling

The early use of the APF has been dominated by two groups, one centered at the University of California at Berkeley (UCB) led by Geoff Marcy with partners at the University of Hawaii, the other centered at the University of California at Santa Cruz (UCSC) led by Steve Vogt with partners at the Carnegie Institution of Washington. Both teams constructed their own nighttime scheduling routines optimized for their own interests, relying heavily on the infrastructure described in previous sections.

6.2.1 UCB Observing Tool

The UCB observing team uses a collection of high-level Python scripts to implement autonomous nighttime operations. The software consistently runs untouched from afternoon calibrations through morning shutdown and data reduction. The level of automation is such that the observer assigned for the night will type a single command at a command prompt, watch the observatory open at sunset, go to sleep, and return in the morning to inspect fully-reduced spectra and/or extracted high-precision radial velocities (RVs) on a web-based display.

First, the spectrograph is focused by taking a series of thorium-argon (ThAr) lamp exposures through the pinhole decker at a variety of dewar focus positions. The FWHM of about 100 thorium lines are measured as

a function of dewar focus position; a polynomial fit is performed to determine the optimal focus, minimizing the FWHM. Once the spectrograph is well-focused a collection of calibration images are acquired, including flat fields, arc lamp exposures, and iodine reference spectra. Two observing plans are then constructed for the night, one for good seeing conditions, and one for poor conditions. Targets are drawn from a shared Google Document, and concatenated into a linear observing script based on target visibility, priority, and RV jitter minimization techniques.¹³

Once the calibrations and afternoon setup are complete the system remains idle until the sun elevation is 3.2 degrees below the horizon. At this time the dome is opened and the telescope is slewed to a bright star to focus the telescope. After the focus is complete it remains on that star to gather data on the sky conditions, which will be used to select from the two available observing programs. Science observations commence when the sun is 8.9 degrees below the horizon, with the scripted observing tool proceeding sequentially through the selected observing plan. While observations are in progress, the higher-level tools continue to assess trends in weather conditions, and the overall progress of the observing plan against its predicted schedule. Clouds are detected by real-time monitoring of the flux received by the exposure meter.¹⁴ If the scripted observing tool is ever more than an hour behind schedule, a new observing plan will be constructed to replace the previous plan. Observations are halted when the sun returns to 8.9 degrees below the horizon.

Data reduction commences shortly after the morning shutdown, with automated creation of a logsheet for the night, and extraction of spectra. The reduced spectra are then passed through a series of calculations to determine the barycentric velocities for all targets and extract the RV shifts relative to an iodine-free reference observation of each target while simultaneously solving for, and removing, the effects spectrograph PSF variability.³ The results are then analyzed by an automated RV analysis code to search for planets and posted to a web-based archive where they can be inspected and interpreted.

6.2.2 Lick-Carnegie Observing Tool and Dynamic Scheduler

The Lick-Carnegie observing team developed a variety of integrated tools to automate nightly operations. Using a single command, the suite runs from afternoon until morning the next day. The suite handles pre-observing exposures, opens the telescope, performs dynamic observations throughout the night, closes the telescope, and ends with post-observing exposures. The sun elevation conventions used by the UCB team were first deployed by the Lick-Carnegie team, with the one departure that the Lick-Carnegie team halts observations when the sun elevation is twelve degrees below the horizon.

While running, the suite checks that sub-tasks are successfully executed, and adapts to changing conditions. Afternoon activities begin with focusing the spectrograph via exposures of iodine lines at a sequence of dewar focus values. The dewar focus is then set to the best-fit value determined by examination of the spectral lines in the iodine region of the detector, optimizing the focus for the symmetry of line profiles. Subsequent calibration frames are taken only after the best focus point is determined. Once completed, activities pause until the sun is sufficiently low to open the dome and mirror cover. On a typical day, that allows roughly one hour for the dome interior to reach thermal equilibrium with the dome exterior before observations begin. After the observing phase is completed, the mirror cover and dome are closed. The suite then initiates morning calibrations, including another pass at focusing the instrument to determine whether the focus shifted during the night.

Currently, the observing phase can be done in as many as three modes. A fixed target list can be specified, or targets can be dynamically selected from a long list of potential targets. The third mode of observing is a mixture of the first two, where targets are dynamically selected after the static list has been exhausted. The procedure to dynamically select a target uses the current telescope pointing, estimated slew times to potential targets, the time remaining until 12-degree dawn, and the current wind speed and direction. It also considers whether a target has already been observed during the night, and if the desired observing cadence has been satisfied; not all targets need to be observed on a nightly basis.

6.3 Future development path

Both scientific teams developed scheduling software to meet their specific needs. To fully exploit the capabilities of the APF facility, a more generalized nightly scheduling script is required. Because the APF is a dedicated facility, with no instrument or focus changes to consider, it has the flexibility to easily schedule observations of

unique events or perform time-domain observations of other objects without significantly impacting the Doppler spectroscopy programs. For example, in the first year of operations, the APF has observed the supernova remnant in M82, observed an RR Lyrae star at the same time as a space observatory and observed an extra-solar planet during transit. Two of these observations required specific timing for the APF, while the supernova follow-up required no special timing but the target had to be inserted into the observation list.

To support a wider array of kinds of targets and scientific objectives, development is underway to establish queued observing at the APF. This new scheduling software would query a series of target priority algorithms, with each algorithm providing a single best target for the current conditions, if any are available, and apply a fairness algorithm on top of the recommendations from each science team. This would allow the addition of targets of opportunity in the middle of the night, such as a recent supernova or gamma-ray burst optical counterpart, ideally without any direct human intervention. In addition, it would be trivial to allow multiple teams to observe in one night, rather than use the classical approach of assigning whole nights to individual teams.

Supporting a single observing queue will require the development of a single nightly scheduler. The intent is to borrow the best ideas of each of the two observing teams' schedulers, and integrate them into a single package; this will also better prepare the APF to support programs from outside of the two core teams.

7. DIAGNOSING UNATTENDED PROBLEMS

One serious operational issue with unattended facilities is determining the nature of problems. With nobody present to recognize and probe unexpected circumstances, it can be very difficult to pin down exactly what circumstances triggered a given problem, especially if the problem is new and unexpected. While verbose logging can help solve a lot of mysteries, the logs still require explicit awareness of a potential problem, as opposed to implicitly containing a complete record of all available data.

The need for always-on information retention was emphasized when the APF began to see regular use under the control of non-technical personnel, who could not reasonably be expected to probe for root causes in quite the same way as would be expected of a seasoned technician. The APF thus became the first major deployment of our KTL keyword history database,⁷ the success of which led to its near-immediate deployment across all facilities at Lick Observatory. For the historian to succeed, the key premise is that all available telemetry is exposed via a single API, and so long as the historian supports that API, all telemetry may be recorded and retained indefinitely.

While there are innumerable beneficial ways to analyze the stored telemetry, the most valuable use has been the analysis of conditions leading up to an unexpected problem with the facility. The ability to look back in time and flesh out the precise circumstances that surrounded a vague report brings an enormous amount of power to the problem solving process. If additional data is required for an analysis, one never has to go back and recreate the circumstances being analyzed, one can instead pull additional data from the existing records and move forward.

In the case of an autonomous, unattended facility, one may not even have a timely problem report to go on. Perhaps an anomaly was noticed in that night's logsheet, or otherwise recognized well after the problem originally occurred. With attended observing, there is at least the hope that the observer may provide a sufficient report to diagnose the issue at hand; with unattended observing, there is no such hope, and the technicians working during the day are limited to analyzing the available data, or attempting to recreate the problem. When put in such a position, there is no substitute for having all the telemetry at your disposal.

8. CONCLUSIONS

Over the past year, the APF successfully transitioned from a commissioning phase to a productive scientific facility capable of radial velocity accuracies better than one meter per second.¹ By the end of the first six months of regular observations, the telescope was observing 75% of the time, including losses to weather. Because of a wet spring in 2014, the total operational efficiency dipped for several months, but only five nights were lost to hardware problems and a comparable number lost to software issues; most of the lost time due to software was a result of ongoing development and testing with the science teams' nighttime schedulers.

Most observations are now routine. Depending on the science requirements and the length of the night, up to 100 science observations are performed on 40-50 individual targets. The regular cadence combined with the high quality of the spectra have already resulted in two planet discovery papers, one on a Neptune mass planet around a M class dwarf¹⁵ and the second on a four planet system found around a K class star.¹⁶

The data from the papers above shows significantly smaller radial velocity errors than Keck/HIRES. Thus, the APF will need fewer observations to discover and characterize planets. The high quality data from the Levy spectrometer also opens up a number of other studies such as characterizing the envelopes of radial pulsating stars, discovering long period binary star systems or monitoring the conditions of other sun-like stars. The ability to rapidly acquire a new target enables follow-up of bright targets of opportunity. Each of these example programs requires regular cadence and high-quality, stable spectrometer. Starting August 1, we will open up 20% of the time on the APF for observers throughout the University of California, and look forward to broadening the scientific impact of this young facility.

The longer term future planet studies will be gathering large samples of planets to answer such questions as how often planets of different masses and radii appear around stars of different types. The APF is well suited for both finding planets for these sorts of population studies and follow-up of planets discovered with both Kepler-II and the Transiting Exoplanet Survey Satellite. In addition, a host of other science topics can be pursued alongside the work-horse planet discovery and characterization programs.

REFERENCES

- [1] Vogt, S. S. et al., “APF - The Lick Observatory Automated Planet Finder,” *PASP* **126**(938), pp. 359–379 (2014).
- [2] Brunswick, R., “Development and testing of an unique telescope enclosure design optimized for seeing and telescope thermal control,” *Proc. SPIE* **5495**, 565–576 (2004).
- [3] Butler, P. R. et al., “Attaining Doppler Precision of 3 M s⁻¹,” *PASP* **108**, 500 (June 1996).
- [4] Vogt, S. S. et al., “HIRES: the High-Resolution Echelle Spectrometer on the Keck 10-m Telescope,” *Proc. SPIE* **2198**(1), 362–375 (1994).
- [5] Radovan, M. V. and Cabak, G. F., “Buckling failure of the Automated Planet Finder spectrometer determine spaceframe,” *Proc. SPIE* **8446**, 84468A–84468A–12 (2012).
- [6] Lanclos, K. et al., “Tuning a 2.4 Meter Telescope... Blindfolded,” *Software and Cyberinfrastructure for Astronomy* **9145**, Proc. SPIE (2014).
- [7] Lanclos, K. and Deich, W. T. S., “A Complete History of Everything,” *Proc. SPIE* **8451** (Sept. 2012).
- [8] Radovan, M. V. et al., “A radial velocity spectrometer for the Automated Planet Finder Telescope at Lick Observatory,” *Proc. SPIE* **7735**, 77354K–77354K–13 (2010).
- [9] Conrad, A. R. and Lupton, W. F., “The Keck Keyword Layer,” *ADASS II* **52**, 203–207 (Jan. 1993).
- [10] Lupton, W. F. and Conrad, A. R., “The Keck Task Library (KTL),” *ADASS II* **52**, 315 (Jan. 1993).
- [11] Grigsby, B. et al., “Remote observing with the Nickel Telescope at Lick Observatory,” *Proc. SPIE* **7016**, 701627–701627–12 (2008).
- [12] Deich, W. T. S., “EMIR: a configurable hierarchical system for event monitoring and incident response,” *Software and Cyberinfrastructure for Astronomy III* **9152**, Proc. SPIE (2014).
- [13] Dumusque, X. et al., “Planetary detection limits taking into account stellar noise. I. Observational strategies to reduce stellar oscillation and granulation effects,” *A&A* **525**, A140 (Jan. 2011).
- [14] Kibrick, R. I. et al., “A comparison of exposure meter systems for three exoplanet-hunting spectrometers: Hamilton, HIRES and APF,” *Proc. SPIE* **6274** (July 2006).
- [15] Burt, J. et al., “The Lick-Carnegie Exoplanet Survey: Gliese 687b: A Neptune-Mass Planet Orbiting a Nearby Red Dwarf,” *ArXiv e-prints* (May 2014).
- [16] Vogt, S. S. et al., “A Four-planet System Orbiting The K0V Star HD 141399,” *ApJ* **787**, 97 (June 2014).