A comparison of exposure meter systems for three exoplanet-hunting spectrometers: Hamilton, HIRES and APF

R. I. Kibrick, D. A. Clarke, W. T. S. Deich, and D. Tucker,

UCO/Lick Observatory, University of California, Santa Cruz, CA 95064 USA

ABSTRACT

The majority of extra-solar planets discovered to date have been found using Doppler-shift measurements obtained with the Hamilton Spectrometer[1] at Lick Observatory and the High Resolution Echelle Spectrometer (HIRES[2]) at Keck Observatory. Each of these spectrometers employs an integral exposure meter which enables observers to optimize exposure times so as to achieve the required signal-to-noise and to determine the photon-weighted midpoint of each science exposure (which is needed to correct the Doppler shift to the Solar System barycenter[*]). In both of these systems, a propeller mirror located behind the spectrometer slit picks off a few percent of the light and directs it to a photo-multiplier tube (PMT) used to measure the exposure level versus time.

In late 2006, the new Automated Planet Finder (APF) Telescope and APF Spectrometer are scheduled to begin operations at Lick Observatory; both will be dedicated exclusively to the search for extra-solar planets. Like the Hamilton and HIRES Spectrometers, the APF Spectrometer will employ an integral exposure meter, but one with a significantly different design.

The APF exposure meter will employ a stationary pellicle located ahead of the slit to pick off 4% of the light and direct it to the guide camera. That camera will produce images typically at a 1 Hz rate, and those images will be used both for autoguiding and for computing the exposure level delivered to the spectrometer. In each guide camera image obtained during a science exposure, the time-tagged signal from the pixels that correspond to the spectrometer slit will be integrated in software to determine the current exposure level and the photon-weighted midpoint of that science exposure.

We compare these two different design approaches, and describe the significant hardware and software features of each of these systems.

Keywords: exposure meter, spectrometer, extra-solar planets

1. INTRODUCTION

The exposure meter systems that are the subject of this paper:

1. Provide in real-time a series of time-tagged measurements of the intensity of light (in a given wavelength range) entering the spectrometer slit
2. Integrate those measurements to determine the current exposure level of the observation in progress
3. Compute the photon-weighted mid-point time for each such observation

---

Further author information- (Send correspondence to R.K.)
R.K.: Email: kibrick@ucolick.org; http://www.ucolick.org/~kibrick; Telephone: 831-459-2262; Fax: 831-459-2298;
D.C.: Email: de@ucolick.org; Telephone: 831-459-2630; Fax 831-459-2298;
W.D.: Email: will@ucolick.org; Telephone: 831-459-3913; Fax: 831-459-2298;
D.T: Email: tucker@ucolick.org; Telephone: 831-459-4903; Fax: 831-459-2298

[*] The barycentric velocity correction includes the following effects, all important at a level about 1 m s^{-1}: velocity vector of observatory relative to the solar system barycenter, time dilation, general relativistic blueshift due to Solar gravitational potential field, modification of stellar coordinates due to proper motion, and the apparent secular acceleration due to the transverse component of the stellar velocity vector.[3] The resulting barycentric corrections are precise to within 0.1 m s^{-1}.

The first function provides observers with useful real-time feedback about short-term variations in the signal level at the slit, which may indicate either variations in atmospheric transparency due to clouds, changes in image quality at the slit due to drifts in telescope focus or changes in atmospheric seeing, or problems with telescope tracking or autoguiding such as may occur during windy conditions. Such feedback can alert the observer to conditions which may require immediate intervention, such as loss of guiding.

The second function enables observers to optimize exposure times so as to achieve the signal-to-noise level required by their science. Such optimization enables observers to make more efficient use of telescope time.

The third function is essential to obtaining Doppler-shift measurements of sufficient precision to support extrasolar planet research programs. At mid-latitudes, the Earth's rotation dominates the time dependence of the barycentric correction over the Earth's orbital motion, resulting in a change of 1.5 m s^{-1} per minute for a typical hour angle and declination. To compute the correction to within 0.1 m s^{-1} (1/10 of the error budget of 1 m s^{-1}), the mid-time of an exposure must be known with an accuracy of 4 seconds.

2. THE HAMILTON SPECTROMETER EXPOSURE METER (1995)

The exposure meter system at the Hamilton Spectrometer evolved from the Coudé integrator. That integrator (also referred to as the exposure integrator or Coudé exposure meter) predates the Hamilton and was originally developed for optimizing the exposure times for observations recorded on photographic plates.[4]

The Coudé integrator monitors throughput at the slit. A reflective chopper behind the slit diverts a small part (approximately 8%) of the beam to a photo-multiplier tube (PMT), and associated electronics integrate the counts from the PMT. The integrator electronics are mounted in a rack in the Coudé slit room, and the integrated counts are displayed on a 7-segment display located in sight of the Coudé observer sitting at the Coudé eyepiece; a dimmer adjustment knob can be used to turn off this display.

But with the advent of the Coudé guide TV, and the move to remote control of observations from a nearby control room, the Coudé integrator signal and the control signal for the dark slide (located in front of the PMT) were extended to that control room (around 1990). A table-top digital counter was used to display the integrated counts and a toggle switch was used to control the dark slide.

In the mid-1990s, the extrasolar planet search program needed to increase the precision of its Doppler shift measurements. In response, the system was upgraded to enable it to calculate the photon-weighted mid-point time for each exposure. Lick Observatory Telescope Technician Wayne Earthman developed the hardware and initial software for an interface between the Coudé integrator electronics and a Compaq PC running MS-DOS.

In 1999, Deich wrote new software on a Linux-based PC and developed an enhanced user interface; he also added software enabling the PC to transmit to the Observatory's data acquisition system the photon-weighted mid-point time computed for each observation, enabling its inclusion in the FITS headers recorded for those observations. Tests indicate that the actual mid-point times of observations are determined to within 1 second.

2.1. Hardware description

A spinning propeller mirror located behind the slit periodically intercepts the diverging beam and directs it to a fold mirror; that mirror in turn sends the light through a small lens and an opened dark slide and onto the PMT. The propeller mirror blade obscures approximately 8% of the area of the annulus it describes when it spins.

The propeller mechanism has a mechanical control, in the form of a plunger on the side of the slit pedestal which, when drawn out, allows it to spin freely, but when put in, stops the propeller in its tracks. When thus stopped by the plunger, the propeller is forced to a position out of the light-path.

The motor that drives the propeller mirror runs at a relatively constant speed, but there is currently no encoding of either the position or velocity of the propeller mechanism, nor of the motor that drives it. Hence, there is no servo control of that mechanism nor any ability to monitor its operation remotely. If there is a problem (e.g., the integrator's counter is not updating), and the propeller is suspected, observers can attempt to determine whether or not it is spinning by entering the Coudé slit room and listening carefully near the slit pedestal. In addition, they can insert the plunger and listen for the characteristic "thunk" of the propeller coming to a halt, and then withdraw the plunger and listen as the motor ramps up.

The housing for the Coudé integrator optics provides for mounting one of two interchangeable PMT assemblies, one blue and one red-sensitive. Exchanging PMT assemblies is a manual operation, but one infrequently performed since the blue assembly is rarely used. The housing also contains an electro-mechanical dark slide (located in front of the PMT assembly) used to protect the PMT from over-exposure; that dark slide is controlled by a toggle switch on the front of the electronics rack (or remotely from the control room).

All of the exposure meter electronics and power supplies are contained in a rack in the slit room, opposite the slit. That rack contains the controls for the integrator electronics, the PMT, and the propeller mirror. It includes: pushbuttons for enabling, disabling, and resetting the counter; a 7-segment display of the integrated counts as they come in; an audible alarm that is triggered if the count rate is excessive; voltage and current meters for the PMT power supply; and the controls and meters used for setting the zero balance and dark current balance for the analog signal processing electronics.

The output from the PMT undergoes analog signal processing that includes circuitry to perform dark current subtraction. The dark-current-subtracted signal undergoes further analog integration and thresholding plus digital prescaling to produce pulses that are counted by a 12-bit counter whose output is used to drive the 7-segment digital readout display.

2.1.1. Interface to the computer

The output pulses from the Coudé integrator electronics are used to toggle a TTL-level signal line connected to a PC's parallel port. A second input on that port is used to monitor a signal that is active when the shutter to the science camera is open.

The pulse-toggle electronics are limited to a maximum rate of 100 Hz, so the PMT trigger threshold is set at a level which keeps the pulse rate well below that maximum, thus ensuring that the pulse rate is linearly proportional to the incident flux rate. Typical pulse rates at the input to the parallel port are < 10 Hz.

The PC is an old Pentium/60 computer with 16 MB of memory, a clock speed of 60 MHz, and the Red Hat 5.0 Linux operating system. It was already an old computer at the time it was installed (late 1999), but it replaced an even older PC that ran MS-DOS.

2.2. Software description

While the software for that older PC computed a photon-weighted exposure midpoint and displayed the result on the PC's screen, it was up to the observer to record those data by hand. The current version broadcasts the exposure midpoint data to another computer on which runs both the graphical user interface (GUI) monitored by the observer and the data-acquisition system, which writes the midpoint data into each image's FITS header.

2.2.1. The midpoint daemon process

The PC runs a small daemon (the midpoint daemon) that monitors the Coudé integrator pulses and science camera shutter signal that are connected to the PC's parallel port. It counts both the positive and negative transitions of any pulses that arrive while the science camera shutter is open, and when that shutter closes at the end of an observation, the daemon computes and broadcasts the photon-weighted exposure midpoint. The computer's internal clock is kept to within a few milliseconds of UTC through the use of Network Time Protocol (NTP) software slaved to a global positioning system (GPS) receiver located at the Observatory.

The midpoint daemon runs in user mode (not kernel mode), sitting in a tight loop and polling the parallel port as rapidly as it can, trying to ensure at least one check of the parallel port every 10 ms. In actual use, this works very well, but in principle the operating system doesn't guarantee that a program will be able to poll the parallel port at any particular rate[†]. Although the midpoint daemon can't guarantee that it samples the parallel port at $>= 100$ Hz, it does check the system time at every sampling. If it detects that over 10 ms. have elapsed since the last check, it increments a "missed sample" counter, and records the missed-sample time in a 1000-element buffer. (If more than 1000 samples are missed, some time-stamps won't be recorded, but in practice that never occurs.)

[†] A better approach would involve generating an interrupt on the parallel port and using a device driver to respond immediately to each interrupt. However, scheduling constraints precluded such an implementation.

When the science camera's shutter is first opened, the midpoint daemon records the absolute start time and broadcasts a MUSIC[5] message indicating that it is beginning to count Coudé integrator pulses for a new exposure. While that shutter remains open, the integrated count and total number of missed samples are broadcast every second. When the shutter closes, the photon-weighted exposure midpoint is computed, and a final message is broadcast containing the start time, midpoint time, end time, and number of missed samples. All of the data from the end-of-exposure message are copied by the data-acquisition system into each image's FITS header.

The complete list of time-stamps of missed samples is recorded to a log file, allowing the diligent observer to compute the uncertainty in the midpoint time resulting from any missed samples. In practice, the number of missed samples is so small (almost always zero) that the broadcast midpoint time is completely satisfactory without needing to consult the log file.

2.2.2. The graphical user interface

A small GUI is provided with the midpoint software (see fig 1); it runs on any observer's computer(s), but never on the midpoint PC itself, so as not to interrupt the midpoint daemon. The GUI is divided into two panels. The upper panel displays a plot of Coudé integrator pulses versus time for the current exposure. The lower panel displays a scrollable history of midpoint data for all exposures, one row per exposure. Each row shows the observation number, start time, midpoint time, end time, total counts, and missed samples.

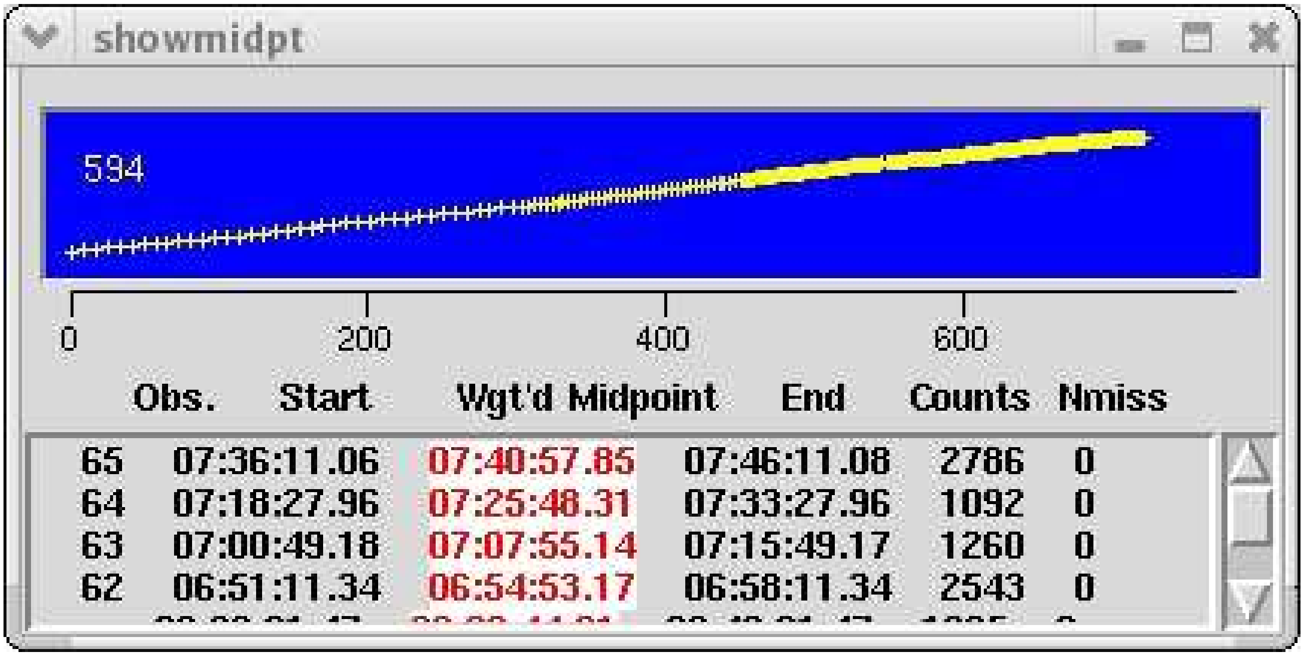

Obs.	Start	Wgt'd Midpoint	End	Counts	Nmiss
65	07:36:11.06	07:40:57.85	07:46:11.08	2786	0
64	07:18:27.96	07:25:48.31	07:33:27.96	1092	0
63	07:00:49.18	07:07:55.14	07:15:49.17	1260	0
62	06:51:11.34	06:54:53.17	06:58:11.34	2543	0

Figure 1. Hamilton exposure meter midpoint GUI

Typically, the observer monitors the integrated exposure level and terminates the exposure when the counts reach the desired level. (It would be simple to provide a user software tool to terminate the exposure automatically when the desired level is reached, but users of the Hamilton Spectrometer have not requested this.)

2.3. Operating procedures

Using the controls on the front panel of the electronics rack, the observer turns on the Coudé integrator electronics (including the high voltage power to the PMT) in the afternoon. After waiting approximately one hour for the electronics to stabilize, the observer can use those front panel controls to set the zero balance for the electronics and calibrate the dark current of the PMT (to enable its subsequent subtraction by the signal processing electronics). Power to the propeller mirror is also switched on during the afternoon and remains on all night (or sometimes through the entire run).

There is a dark slide just in front of the PMT, controlled by a switch on the electronics rack. To protect the PMT from over-exposures, observers are supposed to close the dark slide during arc-lamp and flat-field calibration exposures, and turn the PMT power off in the morning. To date, although accidental over-exposures of the PMT have occurred, none appears to have permanently damaged the device; but such over-exposures do tend to increase the noise level of the PMT signal until the device has had time to recover. In addition, the spectrometer's aperture plate provides some protection in case of the inadvertent use of an over-wide slit.

At the end of the observing run it is very important that the observer always remember to stop the propeller with the plunger before turning the power off. Otherwise, there's an 8% chance that it will coast to a stop in the light path, and the next observer, who may not use the exposure meter, will be mysteriously short on light.

The effective gain of the exposure meter system is determined empirically by comparison to simultaneous observations of standard stars. Those calibrations enable observers to compute what exposure level (as reported by the exposure meter) will result in the desired signal-to-noise ratio for observations of science objects of given magnitudes and spectral types. Since the extrasolar planet search program involves repeated observations of the same objects over periods of years, there is ample opportunity to obtain and refine such calibrations.

2.4. Limitations

The exposure meter for the Hamilton Spectrometer is a simple and robust system whose optomechanical and electronic components have performed well for over thirty years. However, its existing manual controls require that someone (either the on-site observer or Observatory support staff) physically enter the Coudé slit room to perform various functions, i.e., to turn the electronics on and off, to release or park the propeller mirror, or to open and close the dark slide. In addition, during each afternoon of the observing run, that person must enter the slit room to set the zero balance and perform the dark current calibration; that calibration needs to be checked several times during the night, especially in warm weather. Finally, should any problems with its operation arise, the existing system provides little ability to assess its status remotely.

3. THE HIRES SPECTROMETER EXPOSURE METER (2000)

The exposure meter system for the HIRES spectrometer at Keck Observatory evolved directly from the exposure meter system for the Hamilton Spectrograph at Lick Observatory. As a result, it employs a similar design and utilizes similar components, including a spinning propeller mirror (located behind the slit) that briefly diverts the beam to a PMT once a second.

However, in designing the HIRES exposure meter system, a key objective was to overcome many of the limitations of the Hamilton system, particularly the need for observers or observatory staff to enter the slit area to access the manually-operated controls. Every effort is made to eliminate human traffic inside of the HIRES instrument so as to minimize the infiltration of particulate contaminants and avoid elevated humidity levels (from human activity inside the instrument) which can lead to frost developing on slit area components, such as the science camera shutter; frost on that shutter can cause it to stick.

Accordingly, the HIRES exposure meter system (like the rest of the HIRES) was designed to provide remote control of all of its relevant functions, either from an on-site control room, or from the Keck remote observing facilities located at the Keck Headquarters in Waimea, Hawaii,[6] or from similar facilities located in California.[7]

The HIRES exposure meter is one of several upgrades made to the HIRES Spectrometer since it was commissioned on the Keck I Telescope in 1993 . An image rotator was added to HIRES in 1996,[8] the HIRES exposure meter was added in 2000, and in 2004 the original Tektronix/SITE 2K x 2K CCD detector was replaced by a 3 x 1 CCD mosaic comprised of MIT/LL 2K x 4K devices.[9]

3.1. Hardware description

The spinning propeller is driven by a Galil 50-1000 DC servo motor, controlled via a Galil DMC-1510 single axis servo motor controller and Galil AMP-1110 servo amplifier that provide precise velocity control. The use of an encoded servo motor enables remote monitoring of the position and velocity of the propeller mirror. The propeller itself is designed slightly out of balance, such that it coasts to its parked position when its motor power is shut off.

Two Hall effect sensors are attached to the propeller mechanism. One sensor detects when the propeller mirror is in the beam, and is used to trigger PMT samples as the propeller rotates by that sensor. The other sensor detects when the propeller is in the "parked" position, out of the light path. When the propeller is spinning, this "parked" sensor is used to trigger "dark" samples of the PMT, and when the propeller is powered off, this sensor is used to confirm that it is parked (i.e., out of the beam). The signals from these Hall effect sensors are connected to digital inputs on the Galil servo motor controller.

The exposure level is measured with a Hamamatsu HC135-01 PMT sensor module[‡] with a spectral range of 300 to 650 nm and a peak counting efficiency of 22% at 400 nm. Its responsivity is 440,000 CPS/pW, and it provides a dynamic range of 2 x 10^6 with a linearity of ± 1%. This sealed module provides a complete PMT system, including a PMT tube, high voltage power supply, and associated analog and digital processing electronics; it also includes an embedded microcontroller with an RS-232 serial interface that provides a path both for remote control and for acquisition of the data from the PMT. The module is compact (4.75 by 1.375 inches), light weight (180 grams), and requires only a +5 volt supply. The electronics reach stability about 180 seconds after being powered on.

The Hamamatsu PMT is a very delicate device. Normal office light levels can burn out the PMT, if it is powered up. To protect the PMT from high light levels, a photodiode is mounted next to the PMT. The photodiode activates before the ambient light level is high enough to damage the PMT. The signal from the photodiode cuts power to the PMT. A digital input on the Galil servo motor controller monitors the photodiode signal.

The RS-232 serial interfaces from the Hamamatsu PMT and the Galil servo motor controller are connected to a Lantronix ETS 8P terminal server. The terminal server provides remote access from the local area network for the HIRES instrument.

3.2. Software description

The HIRES exposure meter software provides a superset of the functionality provided by the software for the Hamilton system. In addition to providing the basic functions that enable monitoring of the integrated exposure level and the computation of the exposure midpoint time, the HIRES system provides extensive control of and feedback on its various hardware components. This facilitates both remote operation, and remote diagnosis should problems arise. In addition, the HIRES software implements a more accurate method of compensating for thermal drifts of the PMT output.

The stability of the PMT baseline drifts with temperature (11% per degree C) as does the responsivity (±0.1% per degree C). To compensate for these drifts, on each revolution of the propeller mirror the PMT is sampled twice: once, during the time that the mirror is in the beam, as indicated by the "in-beam" Hall effect sensor, and again when it is out of the beam, as indicated by the "parked" Hall effect sensor. In the first case, the PMT is sampling the light coming through the slit, while in the second, it is sampling the PMT baseline and the dark count rate inside the HIRES instrument. The exposure meter software subtracts the second (dark) sample from the first (illuminated) sample to derive the actual exposure level (illuminated minus dark). This method provides continuous subtraction of the PMT baseline so as to cancel out even short term drifts and eliminates the need for repeated manual calibration of the PMT baseline during the night, as is required by Hamilton exposure meter.

The HIRES exposure meter software listens for MUSIC messages broadcast by the HIRES CCD subsystem whenever an observation begins (and the science camera shutter opens) or ends (and that shutter closes). It also listens for messages broadcast when an exposure is paused (indicating shutter closure) and resumed (indicating shutter re-opened).

At the beginning of each science exposure, the exposure meter software clears to zero its measurements for that exposure. During the time the science camera shutter is open for that exposure, it integrates separately the illuminated samples and the dark samples from the PMT as well as the derived signal (illuminated minus dark). In addition, the time-tagged samples of the derived signal are used to compute the photon-weighted midpoint of the exposure.

[‡] http://sales.hamamatsu.com/index.php?id=13199749

The exposure meter control system can also be armed to terminate the science exposure at a preset integrated signal level; this is the mode in which most observers use it. At the end of the exposure, all of the relevant values from the exposure meter system are written as keywords in the FITS header recorded for that exposure.

The HIRES exposure meter software consists of the following components:

1. A GUI that provides interactive control of the system
2. A KTL[10] keyword library that provides a well-defined application programming interface (API) between that GUI and the lower-level control processes.
3. A "dispatcher" process that runs on the instrument control computer and provides the interface to hardware devices (i.e., the Galil motor controller and the Hamamatsu PMT sensor module)
4. Control tasks running on the Galil motor controller that provide for low-level control of the propeller mirror motor and safety interlocks for the PMT sensor module

The GUI is written in Tcl/Tk. The KTL library and the dispatcher are written in C. The control processes on the Galil controller are written in Galil's specialized control language.

3.2.1. Graphical user interface

The user interface consists of two main areas: the PMT/mirror subsystem is shown on the left, and the exposure terminator logic is on the right (see Fig 2 and http://spg.ucolick.org/spie/2006/Hires.Expometer).

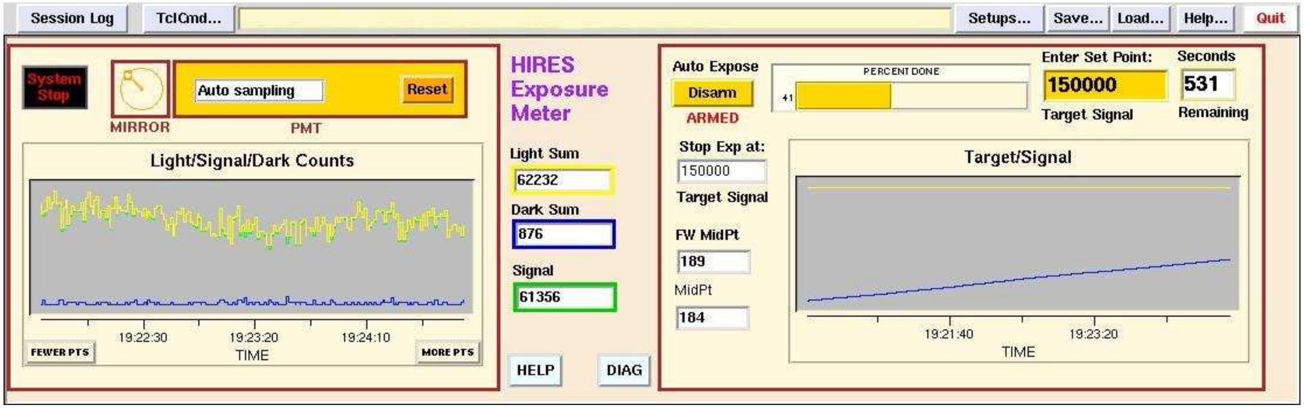

Auto sampling

Light Sum 62232

Dark Sum 876

Signal 61356

Enter Set Point: 150000 Target Signal

Seconds 531 Remaining

Stop Exp at: 150000 Target Signal

FW MidPt 189

MidPt 184

Figure 2. HIRES exposure meter GUI: top-level view

The observer must power up the exposure meter before using it. (When the system is off, the "System Stop" button is replaced by a "System Start" button".) When started, the system takes about 3 minutes to warm up. During that period, a pop-up progress bar provides visual feedback. When the system is ready to use, the mirror will be spinning *and* the status message on the PMT rectangle will say "Ready". (A coarse animation represents the spinning propeller arm).

Once the system is ready, the observer can arm and disarm it with the arming button in the exposure control area. The observer sets the "Target Signal" (number of counts to acquire) by simply typing a value in the "Target Signal" set point box at the upper right. When the system is armed, various color changes will take place and a red "ARMED" message will appear. The GUI was designed (according to Tufte's[11] principles) to provide simple but visually redundant indicators for the various states of the exposure meter control system.

During science exposures, the right-hand plot shows the running sum of acquired signal as a blue line approaching the target signal (yellow line). If the meter is armed, the exposure will be terminated just as the two lines meet. The X axis is labeled in user-readable time (hours, minutes), and an estimated-time-remaining quantity (in seconds) is displayed prominently at the far upper right, so that the observer has a good idea of how long it will take to complete the exposure under current seeing and weather conditions.

Each plot, by default, records about 200 data points. Either plot may be cleared at any time by double-clicking on it; this will discard all plotted data and start over again, and observers may wish to do this if the time scale becomes too compressed.

On the left-hand plot, the user can alter the number of points displayed by using the two small buttons, “MORE PTS” and “FEWER PTS”. The first button will increase the number of points shown on the plot before the plot buffer starts to roll, while the second button will shorten this buffer. The effect is somewhat similar to zooming, but it’s destructive; shortening the buffer and then lengthening it again will discard points clipped during shortening. Usually, observers attempt to choose a plot buffer size that approximately matches the length of one exposure.

The exposure chronological midpoint and the flux-weighted midpoint will be preserved as FITS keywords in the image header. They are expressed in seconds since the start of exposure, and the keywords are EXM0MIDP and EXM0FWMP. They are shown to the left of the “Target/Signal” plot. These values, like the light/dark sums, are updated during exposure, and are cleared at the start of the next exposure.

3.2.2. Service / keyword model

Like all HIRES subsystems, the exposure meter is controlled using KTL keywords.[12] These keywords are contained within a separate exposure meter keyword sharable library (libexpo_keyword.so) and are used in conjunction with the Keck Task Libraries (KTL).[10] These keywords, which are similar in structure to FITS keywords, provide a simple and consistent applications programming interface (API) for interacting with the underlying hardware. For example, the current state of a device can be obtained by reading the keyword corresponding to that device; the device can be commanded to a new state by writing a new value to that keyword. Callback routines can be mapped to specific keywords so that a callback is invoked whenever its associated keyword changes value. These same keywords are also used to document within FITS headers the value of each keyword at the time a given exposure was taken.

Observers generally do not use these keywords directly, since the GUI provides the primary method by which the observer interacts with the exposure meter. However, these keywords can also be used directly from within various types of scripts (e.g., csh, Tcl) to perform automated sequences of commands for either diagnostic or observational purposes.

The HIRES exposure meter KTL keyword library provides three sets of keywords: one set provides access to the propeller mirror mechanism, another to the PMT sensor module, and the third to a virtual device that reports both the instantaneous and integrated values of the dark and illuminated PMT samples and the derived signal.

3.2.3. Dispatchers

The dispatcher process is a device control interface that accepts MUSIC message requests originating from KTL keyword clients, such as the HIRES exposure meter GUI, and translates these into device-specific messages to the attached devices. The MUSIC messages request KTL keyword data or state changes from the devices controlled by the dispatcher.

The dispatcher can control systems consisting of related devices. The dispatcher treats the HIRES exposure meter system as a device with a propeller mechanism and a Hamamatsu PMT sensor module as part of its components. The Hall effect sensor signals and signals from the photodiode are also treated as separate components of the HIRES exposure meter, even though they are accessed via the Galil motor controller.

The dispatcher monitors and controls the status of the PMT, the photodiode (used to power down the PMT in case of overexposure), the propeller motor, and the hall effect sensors. The dispatcher treats all of these individual components as part of an integrated exposure meter device.

The dispatcher communicates with the Hamamatsu PMT sensor module and the Galil servo motor controller via socket connections to specific TCP/IP ports on the Lantronix terminal server. The Lantronix maps those ports to a corresponding set of RS-232 serial ports that connect to the serial interfaces on the PMT sensor module and the motor controller.

In addition to listening for requests from the GUI or other KTL clients (e.g., scripts), the dispatcher listens for HIRES CCD keyword broadcasts of exposure status (exposure integrating, paused, resumed, stopped, aborted). The dispatcher also sends a stop exposure request to the CCD subsystem, to end an exposure after a target exposure level is reached.

3.2.4. Tasks on the Galil controller

Three tasks run on the Galil motor controller. One task handles requests to move the propeller by spinning it at a specified velocity or moving it to specific position, such as the parked position. A second task sends a regular trigger pulse to the PMT when the propeller is stopped or moving very slowly. This enables continuous monitoring of the PMT signal level by the dispatcher. The last task provides an additional safety net for the PMT; it turns off the PMT power when communication with the dispatcher is lost.

3.3. Summary

Observers seem to find the HIRES exposure meter easy to use. It serves both to automate the optimization of exposure time, and as a constant quick-look indicator of seeing and weather conditions. It is easy to see, for example, that high cirrus is drifting into the field: the signal level smoothly declines. This real-time qualitative feedback helps users to abandon unproductive exposures and seek better candidate objects, thus reducing the loss of valuable telescope time.

4. THE APF EXPOSURE METER (2006)

The design for the APF exposure meter is fundamentally different from that of the Hamilton or HIRES exposure meters. Rather than using a propeller mirror located behind the slit to direct light to a PMT, it will instead employ a stationary pellicle located ahead of the slit to pick off 4% of the light and direct it to the guide camera. In each guide camera image generated during a science exposure, the time-tagged signal from the pixels that correspond to the spectrometer slit will be integrated in software to determine the current exposure level and the photon-weighted midpoint of that science exposure.

This approach should provide a number of significant advantages. First, it simplifies the mechanical design and the overall complexity of the system by eliminating a continuously moving part (i.e., the propeller mirror), and with it any vibration or air turbulence that the propeller might generate within the spectrometer. Second, it eliminates the PMT, which is a delicate device that can be damaged by overexposure. Third, by using the guide camera (fed by the pellicle) in place of the PMT, the overall accuracy of the Doppler shift measurements (and not just the exposure midpoint component) may be significantly improved.

4.1. Guiding implications

In both the Hamilton and HIRES spectrometers, the slit jaws are aluminized and tilted by a few degrees to the incoming beam so that the stellar light that spills onto the slit jaws is reflected to the guide camera. Because most of the light from the star is going down the slit, there is no image of the target star on which the autoguider can compute a centroid. Since the field of view of the guide cameras (for the Hamilton, HIRES, and APF spectrometers) is typically less than 1 arc-minute on a side, there is usually little likelihood of a sufficiently bright offset guide star being present elsewhere in the field. Accordingly, the autoguider must employ some sort of quadrant-balancing guider algorithm[13] to guide on the light reflected from the slit jaws. Unfortunately, experience at both Lick and Keck Observatory indicates that such quadrant-based algorithms typically are not as robust and do not perform as well as centroid-based guider algorithms.

By using a pellicle ahead of the slit to pick off 4% of the light, the APF guide camera will have access to the full image of the target star, thus enabling the APF autoguider to employ a centroid-based guiding algorithm. If, as a result, superior guiding is achieved relative to that obtained from guiding using stellar light reflected from the slit jaws, this has the potential to increase the number of photons going through the slit, thereby increasing the photon efficiency of the system and improving the S/N of observations for a given exposure time. If superior guiding also improves the stability of the spectrometer PSF, that could result in an improvement in the overall accuracy of the Doppler shift measurements.

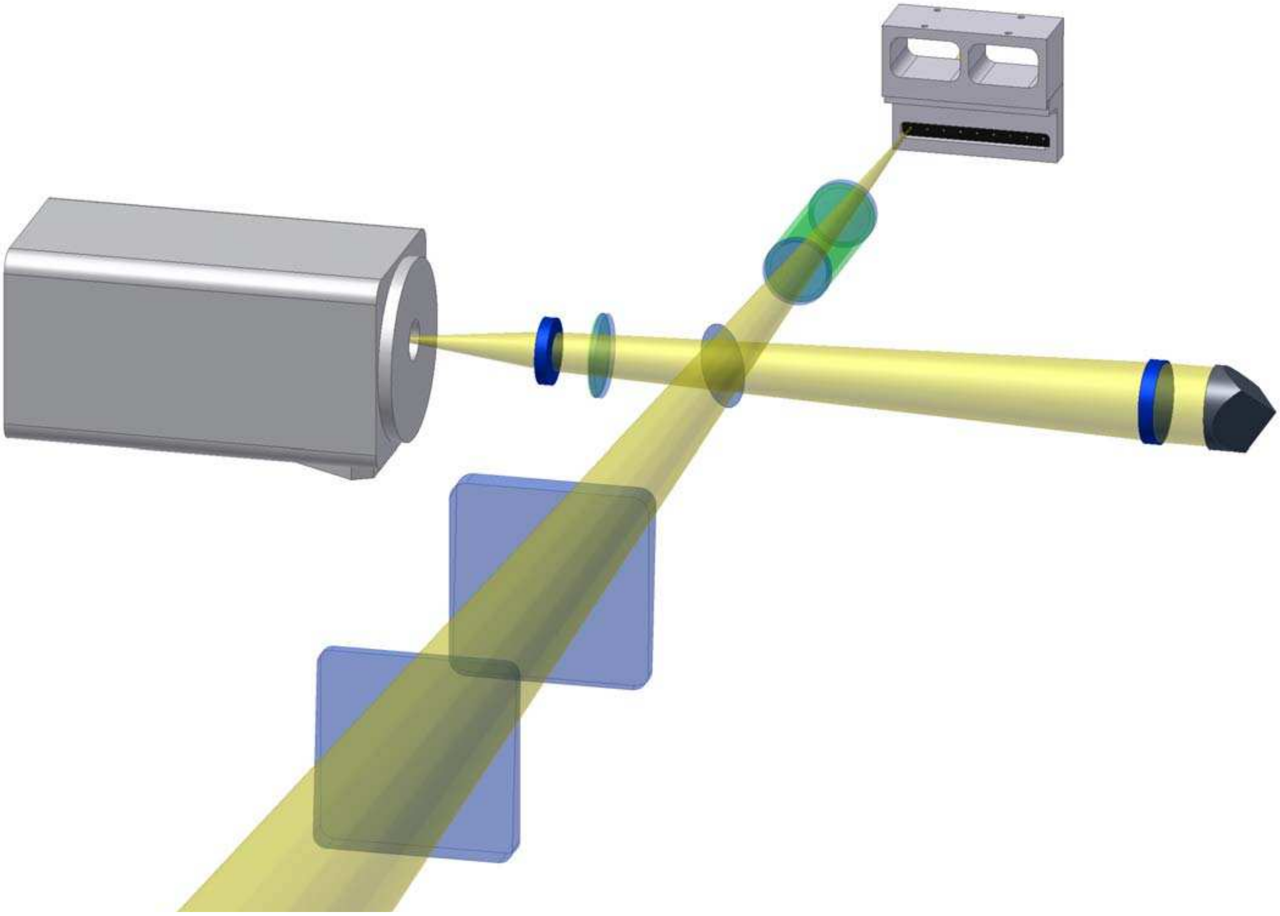

Figure 3. APF exposure meter elements

4.2. Hardware description

Fig. 3 provides a view of the portion of the APF spectrometer where the elements of the APF exposure meter are located. The two main elements are the guide camera (shown at the left side) and the pellicle (shown at the intersection of the two light beams).

Light from the tertiary mirror of the telescope enters at the bottom of the drawing and passes through the two prisms of the atmospheric dispersion compensator (ADC) before hitting the pellicle. Most of the light is transmitted by the pellicle and passes through the iodine cell before converging on one of the discrete slit apertures housed in the aperture select stage in the top right corner of the drawing.

The front surface of the pellicle picks off 4% of the light from the telescope beam and sends it through focal reducing optics to the guide camera (located to the left of the pellicle). To the right of the pellicle is a corner cube that is normally blocked by a shutter (not shown). When that shutter is opened, the guide camera can look through the pellicle to the corner cube which returns an image of the spectrometer slit that is reflected off the back surface of the pellicle. This enables identification of the set of pixels in the guider camera's CCD that correspond to the image of slit.

The CCD camera is a PhotonMax 512B camera made by Princeton Instruments. The camera provides all-metal, hermetic vacuum seals and thermoelectric cooling of the CCD to -80 degrees C. The camera provides digitization to 16-bits at speeds of 10, 5, and 1 MHz. Its interface is a PCI card which can be used with computers running Windows 2000/XP, Linux, or Mac OS X.

The camera utilizes an e2v CCD97 back-illuminated frame-transfer CCD with on-chip multiplication gain. The CCD format is 512 x 512 x 16-μm imaging pixels, yielding an optically-centered imaging area 8.2 x 8.2 mm. With the currently-planned focal reducing optics, the guider will have a field of view about 1 arc-minute square. The CCD97 provides a peak quantum efficiency of 95% between 500 to 600 nm, dropping to about 10% at 300 and 1000 nm. At the nominal operating temperature (-80 C), it has a dark current of 0.05 e-/pix/s.

The CCD97 has two different readout amplifiers (which can be selected under software control), one conventional and the other providing on-chip charge multiplication between 1 and 1000x. The conventional readout amplifier has a typical linear full well of 200 ke-, with read noise of 8 e- rms at 1 MHz and 15 e- rms at 5 MHz. The charge multiplying amplifier provides a linear full well of 800 ke-, with read noise of 45 e- rms at 5 MHz and 60 e- rms at 10 MHz if the on-chip gain is disabled; with on-chip multiplication gain enabled, read noise for this amplifier is effectively reduced to $<$ 1 e- rms.

The availability of this charge multiplying readout amplifier will enable the CCD97 to function more like a photon counting device (similar to the PMTs used in the Hamilton and HIRES exposure meters); thus (even when observing fainter targets) enabling the camera to provide images suitable both for autoguiding and for performing the exposure meter functions.

4.3. Software description

Princeton Instruments provides several layers of software with the PhotonMAX 512-B camera, including the device driver for the PCI card, the PVCAM library (which provides the API) and various user interface and diagnostic software.

The autoguider system for the APF Telescope is being developed by EOS Technologies (EOST), Inc., the contractor responsible for building the telescope. The PhotonMAX 512-B camera and the computer to which it is attached are deliverable components of that system. In addition to developing the autoguider software, EOST will provide to UCO/Lick Observatory (UCO/Lick) a software interface between their autoguider software and the APF exposure meter software being developed by UCO/Lick.

After the telescope acquires a target object and initiates autoguiding, the autoguider software will promptly deliver to the exposure meter software (EMS) the relevant pixels of each subsequent guide camera image and will accurately time-stamp those images to enable the EMS to calculate the photon-weighted mid-point times of science exposures.

By opening the shutter that normally blocks the corner cube, the guide camera can obtain an image of the spectrometer slit, thus enabling identification of the CCD pixels that map that slit image. The guide camera images that are delivered to the EMS will include those slit image pixels and several rings of surrounding pixels.

The pixels that map the slit image will be summed to produce a measure proportional to the light that goes down the spectrometer slit (corresponding to the “illuminated” signal measured by the PMT in the case of the HIRES exposure meter). The EMS will divide into several separate sets the pixels that are outside the vicinity of the slit image; each of those sets will be separately summed to produce a measure of the CCD baseline and sky background, and the sums for each set compared to yield an average baseline/background value exclusive of outliers. That average baseline/background value will be subtracted from the “illuminated” signal to yield the net exposure level for that image. The remainder of the processing of these time-stamped values for the illuminated, baseline/background, and net signal values will be essentially the same as that performed by the existing HIRES exposure meter system as described in Sect. 3.2.

5. CONCLUSION

The exposure meters for the Hamilton and HIRES Spectrometers are an essential component of the exoplanet research programs at Lick and Keck Observatories. Both represent a proven method for providing the functions of monitoring the exposure level and computing the photon-weighted midpoints of exposures.

The exposure meter system for the APF Spectrometer represents a new (and so far unproven) method for providing this functionality. Given the very close proximity of the APF Telescope to the Shane 3-m and CAT Telescopes (both of which can feed light to the Hamilton Spectrograph and its exposure meter), it should be possible to conduct simultaneous observations with the APF and Hamilton Spectrometers and to compare the performance of their respective exposure meter systems. Should it be determined that the guide-camera-based APF exposure meter system does not achieve the requisite level of performance, space has been left within the APF Spectrometer to accommodate a more conventional exposure meter system similar to the one used in HIRES.

ACKNOWLEDGMENTS

We acknowledge and thank the following UCO/Lick Observatory faculty and staff: Dr. Steven S. Vogt, for developing the concept for, and overseeing the development of, both the Hamilton and HIRES spectrometers; Wayne Earthman, for designing and implementing the first implementation of the Hamilton Spectrograph exposure meter system; Tony Misch for providing many of the historical and operational details for the Coudé integrator system; Lee Laiterman, for help in preparing the drawings for this paper; Bernie Walp, for reviewing this paper; and Barry Alcott, Jim Burrous, Ted Cantrall, David Cowley, Erich Horn, Terry Pfister, Jim Ward, Vernon Wallace, and Chris Wright for their various contributions to the HIRES exposure meter system. We also thank Dr. Debra A. Fischer of San Francisco State University and Dr. Geoff Marcy of U.C. Berkeley for their review and comments on various portions of this paper. We appreciate funding from NASA grant NAG5-4445 (1997) to Drs. Vogt, Marcy, and Butler for the design, fabrication, and installation of the HIRES exposure meter and from NASA grant NNG05GK92G to Dr. Marcy for supporting this research. Dr. Fischer is a Cottrell Science Scholar of Research Corporation and acknowledges support from NASA Grant NNG05G164G that made this work possible. We also thank NSF for its grants AST-0307493 and AST-9988358 to Dr. Vogt.

REFERENCES

1. S. S. Vogt, "The Lick Observatory Hamilton Echelle Spectrometer," *PASP* **99**, pp. 1214–1228, 1987.
2. S. S. Vogt *et al.*, "HIRES: The High Resolution Echelle Spectrometer on the Keck Ten-Meter Telescope," in *Instrumentation in Astronomy VIII*, D. L. Crawford, ed., *Proc. SPIE* **2198**, pp. 362–375, 1994.
3. C. McCarthy, "Doppler Corrections to the Solar System Barycenter," Master's thesis, San Francisco State University, 1995.
4. D. Soderblom, *Coudé Spectroscopy at Lick Observatory*, vol. 18 of *Lick Observatory Technical Reports*, UCO/Lick Observatory, 1976.
5. R. Stover, *MUSIC: A Multi-User System for Instrument Control*, vol. 54 of *Lick Observatory Technical Reports*, UCO/Lick Observatory, http://spg.ucolick.org/Docs/KTL/music.ps, 1991.
6. R. Kibrick, A. Conrad, and A. Perala, "Through the Far Looking Glass: Collaborative Remote observing With the W. M. Keck Observatory," *ACM Interactions* **5**(3), pp. 32–39, May-June 1998.
7. R. I. Kibrick, B. Hayes, S. L. Allen, and A. Conrad, "Remote observing with the Keck Telescopes from multiple sites in California," in *Advanced Global Communications Technologies for Astronomy II*, R. I. Kibrick, ed., *Proc. SPIE* **4845**, pp. 80–93, 2002.
8. R. Kibrick, A. Conrad, J. Gathright, and D. Tucker, "User Interface and Control Software for the HIRES Image Rotator on Keck-1," in *Telescope Control Systems II*, H. Lewis, ed., *Proc. SPIE* **3112**, pp. 187–198, 1997.
9. R. I. Kibrick, C. Wright, S. L. Allen, and D. A. Clarke, "Optimization of SDSU-2 CCD controller hardware and software for CCD mosaics," in *Advanced Software, Control, and Communication Systems for Astronomy*, H. Lewis and G. Raffi, eds., *Proc. SPIE* **5496**, pp. 178–189, 2004.
10. W. F. Lupton and A. R. Conrad, "The Keck Task Library (KTL)," in *Astronomical Data Analysis and Software Systems II*, R. J. Hanisch, R. J. V. Brissenden, and J. Barnes, eds., *ASP Conference Series* **52**, pp. 315–320, 1993.
11. E. Tufte, ed., *Envisioning Information*, Graphics Press, Cheshire, CT, 1990.
12. A. R. Conrad and W. F. Lupton, "The Keck Keyword Layer," in *Astronomical Data Analysis and Software Systems II*, R. J. Hanisch, R. J. V. Brissenden, and J. Barnes, eds., *ASP Conference Series* **52**, pp. 203–207, 1993.
13. R. Kibrick and L. Robinson, "The Lick Observatory TV Autoguider," *PASP* **99**, pp. 1014–1021, September 1987.