

HIPO data products

E. W. Dunham^a, J. L. Elliot^{a,b}, T. A. Bida^a, P. L. Collins^a, B. W. Taylor^c, and S. Zoonematkermani^a

^aLowell Observatory, 1400 W. Mars Hill Road, Flagstaff, AZ 86001

^bDepartment of Earth, Atmospheric and Planetary Science, MIT 54-422, 77 Massachusetts Ave.,
Cambridge, MA 02139

^cBoston University, IAR, 725 Commonwealth Avenue, Boston, MA 02215

ABSTRACT

HIPO is a special purpose instrument for SOFIA, the Stratospheric Observatory For Infrared Astronomy. It is a high-speed, imaging photometer that will be used for a variety of time-resolved precise photometry observations, including stellar occultations by solar system objects and transits by extrasolar planets. HIPO will also be used during the test program for the SOFIA telescope, a process that began with a series of ground-based tests in 2004. The HIPO requirements, optical design, overall description, and an early look at performance and planned data acquisition modes have appeared in earlier papers (e.g. Dunham, et al., Proc. SPIE 5492, 592-603 (2004)). This paper provides an update to the instrument description, final lab measurements of instrument performance, and a discussion of the data produced by the various observing modes.

Keywords: CCD, airborne, imaging, photometer, SOFIA

1. INTRODUCTION

HIPO is a first-light instrument for SOFIA that has been designed for observations of stellar occultations¹ and transits by extrasolar planets², and for performance testing of the SOFIA telescope. A general overview of the instrument³ and its optical design⁴ have been published in earlier papers. The present paper updates the instrument description (Section 2) and describes its data products in some detail (Section 3). Section 4 is a very brief discussion of near-term additions to the capabilities of HIPO.

HIPO possesses a variety of time-resolved CCD readout modes. Some of these are needed for occultation and transit observations, but telescope performance testing motivates others, particularly the specialized fast modes. Image motion due to telescope pointing and control system problems and vibrational modes in the telescope are likely to be of concern. The lowest frequency modes in the telescope are at ~25 Hz with significant modal amplitudes ranging up to ~100 Hz. HIPO has a range of observational modes that can sample these frequencies. The highest speed effect we anticipate seeing in SOFIA data is the speckle pattern from the turbulent shear layer. The characteristic time for this is the time required for turbulent structures in the shear layer to move one turbulent scale length at roughly half the speed of the aircraft. This is approximately 500 microseconds, requiring a frame rate of ~2KHz to resolve.

2. INSTRUMENT UPDATE

Since our previous paper³ (hereafter referred to as Paper I), the instrument has been used for ground-based tests of the SOFIA telescope, its hardware has been completed, it passed its FAA “conformity” inspection, and all of its readout modes are functional. The high-level control software is evolving to a two-level configuration with a simplified LOIS⁵ control layer communicating with a graphical user interface (GUI) based on the Eclipse Rich Client Platform (http://wiki.eclipse.org/index.php/Rich_Client_Platform) across the network using the Apache ActiveMQ message broker (<http://activemq.apache.org>). We were prompted to make this change in part to implement our planned network-based control software architecture and in part by instability in the Tk-based GUI in LOIS. LOIS retains its modular character and Tcl-based scripting capability.

The completed instrument is shown in Figure 1 and the as-built instrument performance is summarized in Table 1, an updated version of the performance table first shown in Paper I. We discuss below the areas of the table that have changed since Paper I. In other respects the discussion presented there remains valid.

Table 1. HIPO performance requirements.

Detector	Verification Method	Requirement	Actual
Read noise, 1 Mpx/s	Measurement	6 electrons rms	8 e ⁻ at 1 Mpx/s 5 e ⁻ at 666 Kpx/s
Read noise, 20 Kpx/s	Measurement	3 electrons rms	3 e ⁻ for ≤ 200 Kpx/s
Peak quantum efficiency	Vendor measurement	≥80%	82% (blue), 88% (red)
0.35-0.85 micron QE	Vendor measurement	≥40%	42% minimum
Full frame read rate	Demonstration	≥2 Hz (unbinned)	2 Hz (see text)
Three 80x80 subframes	Demonstration	≥50 Hz (3x3 binned)	11.5-22.2 Hz (see text)
One optimal 80x80	Demonstration	≥100 Hz (3x3 binned)	80 Hz (see text)

Filters	Verification Method	Requirement	Actual
Johnson UBVRI	Vendor measurement	Standard bandpasses	Standard bandpasses
Methane filter, 0.89 μm	Vendor measurement	892 nm λ _c , 17-19 nm δλ ≥80% T _{peak} , 10 ⁻⁵ blocking	892 nm λ _c , 17.8 nm δλ 90% T _{peak} , 10 ⁻⁵ blocking

Optics Configuration	Verification Method	Requirement	Actual
Number of channels	Design	Two	Two
Differential Alignment	Measurement	≤ 100 pixels and 1°	40 pixels, < 0.5° (Typ)
Pupil viewing optics	Design	Red, blue, & high res.	Red, blue, & high res.
Shack-Hartmann optics	Design	Red channel only	Red channel, 3 lenslets
Retroreflection optics	Design	Present	Present
Flange fiducial	Measurement	Centered to ±0.1 mm	Centered to < ±0.05 mm
Bare CCD, 0.055"/px	Design	Present	Present, w/ filter wheel

Optical Performance	Verification Method	Requirement	Actual
Unbinned image scale	Measurement/analysis	1/3"/pixel	0.327 "/pixel (red) 0.331 "/pixel (blue)
3x3 binned image scale	Measurement/analysis	1"/pixel	0.981 "/pixel (red) 0.993 "/pixel (blue)
Field of View	Measurement/analysis	5.6'x5.6'	5.58' (red), 5.66' (blue)
Image quality	Measurement/analysis	80% encl. light in 2x2 px	80% ELD < 2.5 pixel
Distortion	Measurement	≤0.1% (red channel), ≤0.3% (blue channel)	<0.25% (Red; see text) < 0.5% (Blue; see text)
Optics throughput	Measurement/analysis	≥70%, 0.4-0.9 μm	Blue: 66-73% .34-.64μ Red: 72-82% .46-1.0μ

Detector Thermal	Verification Method	Requirement	Actual
Detector temperature	Measurement	-100 to -130C	Settable in this range
Detector temp. stability	Measurement	±0.5C	±0.2C
Cryostat hold time	Measurement	24 hours	25 hrs (lab, heater off)
Cryostat cooldown time	Measurement	NA	6 hours (lab)

Time and Position	Verification Method	Requirement	Actual
Time accuracy	GPS spec & Test	50 μs	3.5 μs late (see text)
Position accuracy	GPS vendor spec	30 meters	5 meters (SA off)

Chopper Control	Verification Method	Requirement	Actual
Chop Trigger Timing	Demonstration	Known to < 100 μs	±1 μs
Chop Outputs	Demonstration	TA_SI_04, §4.2.8	Complies

Table 1. HIPO performance requirements, continued.

Stiffness (HIPO alone)	Verification Method	Requirement	Actual
Flexure	FEA	$\leq 0.1''$ referred to sky	$\sim 0.05''$
Lowest resonance Mirror mounts Detector mounts Optics box	Accelerometer Spring const. & mass FEA	≥ 100 Hz	96, 99, 170, 260 Hz 130 Hz 110 Hz
Sensitivity	Verification Method	Requirement	Actual
Red Channel	Standard Star	NA	See text
Blue Channel	Standard Star	NA	See text
Shack-Hartmann	Lab and sky tests	± 1 pixel spot position	< 0.1 px, see discussion

The read noise has been optimized and the readout modes completed, allowing the noise and read rate requirements to be tested. HIPO nearly meets its high-speed noise requirement and meets its low-speed requirement at a pixel rate an order of magnitude faster than the required rate. The frame rate performance for full frame is met, the single optimal subframe requirement is nearly met (using pipelined occultation mode), but the three-subframe requirement is not. In retrospect, the three-subframe requirement we set for ourselves over 10 years ago was overly optimistic.

During our work aboard SOFIA in 2004 it became clear that precise knowledge of the CCD orientation and location of the center of the instrument mounting flange was needed. We have addressed this need by fabricating a precisely machined fixture called the flange fiducial with a hole pattern that can be imaged by the red channel CCD. Distortion in the red side optics was measured using this fixture, but is small enough that irregularity in the hole pattern is comparable to the distortion. A similar situation pertains for the measurement of distortion in the blue optics.

A new cold strap was installed subsequent to Paper I, bringing the CCD operating temperature to its correct value, but necessarily reducing the hold time of the cryostat somewhat.




Fig. 1. The instrument, complete except for final cable routing, is shown mounted on the SOFIA telescope during ground operations in 2004 (left) and the completed instrument control rack is shown in the lab (right). The laptop computer on the keyboard tray is not part of the rack configuration, but demonstrates that the tray can carry this load.

The timing accuracy specification given in Paper I was only the specification for the GPS clock unit. The performance given in Table 1 is the performance of the full system measured from the 1 Hz output of the GPS clock to the onset of the parallel clocking of the CCD. The observed delay is a combination of hardware propagation delays and software

delays, but is dominated by the delay in the DSP code loop that checks for the CCD readout trigger signal. We have also added a synchronized signal for controlling the SOFIA chopping secondary mirror for test purposes.

The instrument throughput was measured in a series of laboratory and telescope tests using spectrophotometric standard stars. The throughput of the optics is well established by two laboratory tests separated by several years. The telescope tests were problematic in spite of careful atmospheric extinction and telescope reflectivity measurements. The ratio of observed standard star signal to predicted signal was 75% for the V filter and 62% for the I filter with fairly fresh telescope mirror coatings. With old coatings in poor condition these ratios were 63% and 54% respectively. We ultimately concluded that this discrepancy was due to the scattering behavior of the telescope optics and the details of how the reflectometer treated scattering from the reflective surface of the telescope mirrors.

The Shack-Hartmann test sensitivity is dependent on the position uncertainty of the individual spots in the test image. This is related to the signal level and the amount of image blur and motion in the image. In well-exposed lab images the position can be determined to a much higher precision than necessary. In real images of stars with SOFIA on the ground in 2004 the Zernike coefficient reproducibility is generally on the order of 1/10 wave or better. One exception was for astigmatism in the diamond-turned backup aluminum secondary mirror, which was reproducible to about 1/3 wave out of a total contribution of almost 6 waves, corresponding to reproducibility of 5% for this huge aberration component.

The HIPO instrument control rack, shown in the right panel of Figure 1, contains the control computers (including a hot spare) at the bottom. A power control unit and patch panel with connections to the Lights-Off-Management ports of the computers are in the area above the computers on the left side. The corresponding area on the right side is unused. Above this are two keyboard trays, GPS clocks, and flat panel displays. The displays, keyboard, and mouse are stowed and the keyboard trays fold up and are latched in place during takeoff and landing. The display mounting hardware and the keyboard trays were custom built because of the FAA stress analysis and certified materials requirements. Not visible in this figure are two network media converters, two TTL/fiber converters for the timing signals for the instrument, the HIPO network switch, and a spare network switch.

3. DATA PRODUCTS

3.1 Introduction

The HIPO CCDs are operated by Gen II ARC (Leach) controllers⁶ with heavily modified DSP code in the PCI interface and timing boards. The ARC-supplied PCI DSP code was not robust against competing bus activity, such as disk or network data transfers. This was a fatal problem for HIPO since we must write to disk while taking data at the maximum rate possible with Gen II controllers for extended periods of time. We found that the programmed I/O approach taken in the Gen II PCI DSP code was busy about 80% of the time in the absence of competing bus activity. With competing activity PCI retries occurred and the busy fraction rose until data loss occurred and the PCI card hung. This problem was fixed by importing and modifying the DMA code supplied by ARC with their Gen III systems. The busy fraction in the absence of competing activity dropped to about 20% and data transfers are now robust against very severe bus contention. A related modification in LOIS was made in which a static buffer for CCD data suitable for single frames data was replaced by a dynamic circular buffer allowing continuous CCD reads and disk writes. The modified timing board DSP code is complex and its functions are described in the remainder of this section.

Paper I contains a short description of each of the HIPO CCD readout modes planned in 2004. At that time only single frames mode was implemented with HIPO, though slow dots and strip scanning had been implemented on other Lowell instruments. Since then all of the modes except the Series mode have been implemented in the HIPO software, and a new mode, pipelined occultation, has been implemented as well. The Series mode is useful with full-frame CCD imagers, but with a frame transfer CCD like the e2v CCD47-20 used in HIPO, the Basic Occultation mode is a superior way to achieve the same result. We have therefore intentionally deleted the Series mode from the HIPO software.

All of the modes described here can make use of pixel binning with arbitrary binning factors. The DSP code provided by ARC can handle binning in a very general manner, but the loop and subroutine call overhead is unacceptably high for the demands of HIPO. We have augmented this binning code with in-line serial read code for serial binning factors up to 5. Larger binning factors use the general code from ARC with its associated timing penalty. The in-line code segments are stored in high Y memory and copied into fast low memory for use.

Subframing is also supported, in some cases including multiple subframes using the ARC-supplied approach. This allows up to 9 subframes, though we don't plan to use more than three. The ARC approach constrains the subframes to

be the same size, to have no overlap in the vertical direction, and to use a single amplifier readout. We may choose to relax some of these constraints in the future but at present we accept these restrictions. Subframes never include prescan pixels but optionally include postscan pixels. (Full frames do include prescan and optionally include postscan pixels.) There is a substantial timing benefit if postscan pixels are not used. The pixels in the serial register following the subframe are not clocked out – the serial register is cleared using the dump gate in the CCD47-20. Overclocked rows are supported only with full frames in find, single frames, and basic occultation modes. Multiple subframes are not supported in some modes because they don't make sense. These cases are discussed in the mode descriptions below. When multiple subframes are taken the data for each subframe is stored in a separate file.

In all modes except find mode the images can be hardware or software triggered. Hardware triggering causes the image to have a precisely known start time since it is based on a hardware signal derived from the GPS clock. Software triggering has no such timing pedigree and is available as a fallback operating mode that may be useful if care is taken to calibrate the timing or if precise timing is not required. In a time-resolved mode hardware triggering causes the interval between frames (the period of the imaging process) to be defined by the trigger interval and the exposure time is this interval minus the fixed amount of time required for the clocking function involved in that particular mode, subframe geometry, and clock speed. Software triggering, in contrast, uses the exposure timer in the controller so the exposure time is controlled and the interval between frames is the exposure time plus the fixed time required for clocking.

HIPO's hardware trigger signal is ultimately generated by counting 1 MHz pulses from its GPS clock units. The trigger period is defined by a 16-bit counter which can count the 1 MHz signal directly or prescaled pulse trains having 1 KHz or 1 Hz frequencies. As a result it is possible to define trigger periods up to 65.535 milliseconds at microsecond resolution, to 65.535 seconds at millisecond resolution, and to 65535 seconds at 1 second resolution. On the 1 second and 1 millisecond scales the pulse is high for one second and one millisecond, respectively. On the 1 microsecond scale the pulse is high for 10 microseconds owing to a bandwidth limitation of an optical isolator used in the trigger circuit. This has no effect on timing since the rising edge of the trigger is the important one (see Section 3.4 below).

The primary function of HIPO is time-resolved imaging so most of its files are 3-D FITS files with NAXIS3 being the time axis. Single frames and find mode (if the final frame is stored) are the only 2-D HIPO file types. The Strip-scanning mode is an unusual case in that the row dimension is often both spatial and temporal in nature. There are cases where it is purely temporal, however, such as a lunar occultation or other time-resolved 1-D data product. As a result we have elected to treat strip scans as 3-D files in which NAXIS2 is vestigial, having a value of 1.

With the upgraded PCI DSP code and the circular buffer in LOIS the limiting factor on file size is the pixel count approach used in the timing board and PCI DSP code. The timing board sends two 24-bit numbers to the PCI board which multiplies them to form a 48-bit pixel count. For HIPO the first 24-bit number is the product of the number of rows and columns and the second is the product of the number of subframes and the number of images in the time series. This allows files that may in practice exceed a 32-bit operating system file size limit, requiring 64-bit file size compile flags to be used.

3.2 Find mode

In Find mode the shutter is opened and the CCD is read out repeatedly until stopped by the observer. This mode is intended as an aid in field acquisition; as such it is always software triggered and will often be used with binning. Find mode can support multiple subframes, though this is not normally useful. Generally images obtained in find mode are not stored, although the last image may be manually saved as a 2-D FITS file. Find mode uses the same timing board DSP code as single frames with the exception that shutter operation is skipped and handled at the LOIS level.

3.3 Single frames

Single Frames is a simple imaging mode that takes a picture and stores it to disk, incrementing the file index number. It is also possible to request that several frames be taken with one command. The images are straightforward 2-D FITS files. Single frames supports full frames and single or multiple subframes as described above. Figure 2 shows two examples of single frames taken from our work on SOFIA in 2004 illustrating the Shack-Hartmann and retroreflection capabilities of HIPO. The Shack-Hartmann image (left in Figure 2) shows the strong aberrations present in SOFIA's backup diamond-turned aluminum secondary mirror that was used during this period. (A decision was made to do these tests with the backup mirror to avoid risk to the high quality SiC secondary mirror.) This serves to illustrate the way the Shack-Hartmann images work and to show that the backup mirror is suitable for its intended purpose: allowing far-IR observations to be carried out unimpeded in the event that the SiC mirror is damaged and needs to be replaced. The

image on the right in Figure 2 was taken with the retroreflection LED turned on and the spherical button mirror installed on the SOFIA secondary mirror. The mechanical shutter was left open during readout and the LED image left a trail on the CCD during the frame transfer. There was an oscillation of the chopping mechanism's control system at ~ 300 Hz that can clearly be seen in this trail.




Fig. 2. Examples of Single Frames. Left: The contour plots on the left correspond to the pre-alignment condition of the telescope (top) and post-alignment condition (bottom). They are contour plots of the left and right image in the top row of the four-image panel on the right side of the figure. The lower images in the four-image panel are Shack-Hartmann images using the medium density lenslet module for the pre-alignment (left) and post-alignment (right) conditions. Coma was effectively eliminated during alignment. Right: These images were taken with the mechanical shutter in HIPO turned off (i.e. the shutter was open). The streaks above and below the bright image were formed during the frame transfer and pre-exposure flush respectively. The left image of the three shows the case where the chopping secondary system was turned off. The right image has the control system on but set for zero amplitude. The center image shows the system turned on with a 0.1 arcsecond amplitude. The pattern noise visible in these images has been eliminated since these data were obtained in 2004.

3.4 Basic occultation mode

Basic Occultation mode is the most general, but slowest, time resolved imaging mode for HIPO. It uses the same timing board DSP code as Single Frames, and therefore supports full frames and single or multiple subframes as described above. The mechanical shutter remains open during the entire time series observation with the frame transfers defining the integration boundaries. When hardware triggered the shutter opens and the first integration begins on the first falling trigger edge. The first frame transfer is initiated on the rising edge of the trigger signal. The second integration begins when the frame transfer is complete, and is overlapped with readout of the frame store area. The same timing applies to all subsequent integrations. The duration of the first integration is generally different from successive integrations, but at the millisecond level - the frame transfer and shutter opening times are comparable. When software triggered there is not a fixed start time and the period of the time series is the exposure time supplied to the controller plus the duration of the frame transfer and readout. This mode is distinguished by the fact that the full 1033 row frame transfer is done at the completion of each integration. There is no preferred location for the object of interest, although faster rates are possible with subframes closer to the readout amplifier. The maximum frame rate depends on readout details (subframe size and location, binning factor, parallel clock speed, and per-pixel integration time). The images are 3-D FITS files organized as a stack of 2-D images. The third axis is time. These files can become very large.

3.5 Fast occultation mode

Fast Occultation mode is faster but more specialized than the basic occultation mode. It uses distinct timing board DSP code, requires a single subframe, and does not support full frames or multiple subframes. Furthermore, the single subframe is always located with its first row at the bottom of the imaging area of the CCD. Its speed is greatly improved if the subframe is at low column numbers and postscan pixels are not requested. The mechanical shutter remains open during the entire time series observation with partial frame transfers defining the integration boundaries. This mode is distinguished by the fact that only a partial frame transfer (of as many rows as the subframe height) is done at the completion of each integration. At this time the image area begins integrating and the charge image from the previous integration is shifted the rest of the way down to the serial register as the next integration is proceeding. This reduces the dead time due to parallel shifting compared to the basic occultation mode in the ratio of the subframe height to the 1033 rows in the storage area. As with basic occultation the maximum frame rate depends on readout details. The timing details for hardware and software triggers are similar to those described for basic occultation except the duration of the partial frame transfer is shorter than for the full frame transfer so the duty cycle is improved. Because only a partial frame transfer occurs, each image is exposed to sky for a longer period than the integration time on the object of interest. This is normally not a serious issue since this mode will only be used at high frame rates. If necessary a focal plane mask can be introduced to reduce this problem.

Figure 3 shows an example of a fast occultation observation of a test target in HIPO based on a 7-segment LED display driven by a counter counting a synchronized 10 Hz clock. The integration interval for this image was 50 ms, so the LED pattern is the same for pairs of frames. The first state of the 7-segment display is all off, so there should be two blank images at the beginning of the file. In fact there are four. This is because the display was not located immediately above the seam between the image and storage area, but two subframe heights farther up. This serves as a cautionary tale – if a mask is not used and the object is not placed in the first subframe height above the seam between the image and storage areas the data will appear to have a timing shift. Furthermore, if there are additional objects of significant brightness in the same columns as the target object they will appear, with timing shifts, in the frames with the target object. The timing is dependent on the location of the image in a deterministic manner, but the observer must pay attention! These files can become very large.



Fig. 3. Fast occultation mode example. Time increases to the right in this series of images of a test target based on a 7-segment LED display that changes every 100 ms. The frame interval is 50 ms so the target remains the same for pairs of frames. The first state of the target is with all segments off, so there should be two blank images at the left. The fact that there are four is related to the position of the 7-segment display in the image. See text for further discussion.

3.6 Pipelined occultation mode

Pipelined occultation mode is a modification of fast occultation mode that is faster but even more specialized. It uses the same DSP code as fast occultation mode with the exception that the final shift of the image down the frame store region to the serial register prior to readout does not occur. In this case the partial frame transfer of fast occultation mode is followed directly by a read of the same number of rows. The difference between fast occultation and basic occultation is one of duty cycle; the shortest integration interval remains the same since the charge image is shifted the full height of the frame store area prior to readout. Pipelined occultation has the same duty cycle as fast occultation but reduces the shortest possible interval because the bulk of the frame transfer does not occur for each image.

Pipelined occultation mode is tricky to use. Ideally the height of the subframe would be an odd integer submultiple of the height of the frame store area, but that is 1033 rows, a prime number! Fortunately being close is good enough if one is willing to accept a few bad rows at either the top or bottom of the subframe. As an example to see how this works, imagine an unbinned subframe 344 rows tall. Three such subframes fit in the storage area (with one additional row left over). An object within the first 344 rows above the seam between the image and frame store areas leaves a charge image that is shifted down into the top third of the frame store area on the first trigger. After this shift the bottom third of the frame store area is read out, a spurious frame at the beginning of the pipeline. When the read is finished the

original charge image is in the middle third of the frame store area. On the next trigger the second charge image is moved to the top third of the frame store area and the first image is now in the bottom third. When the read happens, the second frame is not spurious, but is the first charge image. If there are $2n+1$ subframe heights in the frame store area there will be n spurious pipeline frames before the first real image. If the number of subframes in the frame store area is even, all of the good images will be dumped during the partial frame transfer and all the reads will be empty.

Pipelined occultation is quite restricted in its application, but it is substantially faster than fast occultation. For example, when measuring the fastest frame rate for an optimally placed 3×3 binned 80×80 subframe for Table 1, the best that fast occultation could do was 50 Hz. By massaging the aspect ratio of the subframe so its height was nearly an odd submultiple of 1033 (69 binned pixels) and expanding the width (96 binned pixels) to retain approximately the same total number of pixels in the subframe, pipelined occultation could reach 80 Hz. This makes the difference between being able to resolve the “rocking modes” in the telescope or not.

3.7 Fast dots mode

Fast Dots mode^{7,8} was developed over 20 years ago to solve the problem of freezing the shear layer speckle pattern seen on the Kuiper Airborne Observatory. It is the fastest mode in the HIPO arsenal, and uses its own unique timing board DSP code. In this mode a finite string of images can be stored at very high frame rates using the CCD as an analog charge image storage buffer. The star image is placed on the CCD near the edge farthest from the serial register to maximize the available analog storage area. The mechanical shutter is opened on the first falling trigger edge, and on every rising trigger edge the parallel clocks (both image and frame store clocks) shift the integrated charge image toward the serial register by a specified number of rows. This timing is analogous to the three occultation modes. The case for software triggering is similar as well – the interval between images is the requested exposure time plus the parallel clocking time. However, the exposure timer in the CCD controller works on a 1 KHz timebase so high speeds can only be obtained with hardware triggers. When the requested number of shifts has been carried out the shutter is closed and the CCD is read out. Streaks are frequently seen between the images in a fast dots frame, being more pronounced at higher frame rates. The integration interval is limited only by the number of rows per shift specified and the parallel clock speed. The width of the subframe specified does not affect the data acquisition speed, only the readout time of the completed image. Fast dots can use the full width of the CCD or a single subframe, but not multiple subframes. If a single subframe is used, only the horizontal placement and size of the subframe are used.

If the file is displayed as a 2-D image it appears to be a string of images (dots) arranged vertically and spaced, on average, by the number of rows shifted between integrations. However, for consistency with the other HIPO time-resolved readout modes we have chosen to represent a fast dots dataset as a 3-D FITS file. It will often be the case for images obtained in this mode that the mechanical shutter's opening and closing times will be very slow compared to the integration interval. This may cause the first images to be faint or entirely missing while the last image may be excessively bright. It may be necessary to request substantially more images than the CCD can actually hold in order to avoid completing the entire operation before the shutter has opened. The timing of the images obtained under these conditions must take into account the number of images lost during the data acquisition cycle.

3.8 Slow dots mode

Slow Dots mode uses its own distinct DSP code, but resembles both fast dots and pipelined occultation modes. It is similar to fast dots mode but instead of separating the individual images by parallel shifts only, the images are separated by full readout of the number of rows in the subframe. Slow dots is also similar to pipelined occultation mode except that the partial frame transfer is not done. It therefore has a constraint similar to pipelined occultation in that the subframe height should be nearly an integer submultiple of the frame store area's height, 1033 rows, and the object of interest should be located in the first subframe height above the seam. Slow dots mode can produce an arbitrarily large number of images at the expense of a much longer shifting period, which is clearly dependent on the width of the subframe. To avoid smearing between images the shutter is actuated on each image and is closed during the read periods. Slow dots can use the full width of the CCD or a single subframe width, but not multiple subframes. When hardware triggering is used the shutter opens and the first integration begins on the first falling trigger edge. Readouts are initiated on each rising trigger edge. The integration time is the trigger period minus the readout time. Software triggering works as described in the basic occultation mode section.

The data produced by this mode is again a 3-D FITS file. This file will include as many blank frames at the beginning as there are subframe heights in the 1033 rows of the storage area of the CCD. These blank frames are removed prior to writing the data file. The caution described in the fast occultation section regarding timing offsets that occur if the object

is not in the first subframe height above the seam applies to slow dots as well. Generally speaking, in a frame transfer CCD the pipelined occultation mode is a better way to achieve the results of slow dots provided that the subframe height constraint for pipelined occultation is acceptable. The slow dots mode is more useful with full frame CCDs where the various occultation modes are not possible due to the lack of split parallel clocks.

3.9 Strip-scanning (TDI) mode

Strip-scanning (or time-delay and integrate, TDI) mode uses the slow dots mode DSP code with a one-row subframe height and with the shutter opening at the beginning of the strip and closing at the end rather than opening and closing for each row. Unlike slow dots mode, no rows are deleted prior to writing the data file. This mode has been used for a number of applications over the years and may be useful on SOFIA for testing the telescope scanning operation with the scan direction parallel to the CCD columns. Also, if the image is binned by a large factor in the parallel direction it is possible to obtain a high-speed one dimensional star position over long time periods. Finally, if a coarse grism is incorporated in the HIPO optical system along with an aperture at the focal plane, a spectrally resolved time series can be obtained at high speed. This would be appropriate for observing lunar occultations.

4. IMPENDING ADDITIONS TO HIPO

We currently have two e2v CCD67 CCDs on order for optional use in HIPO. These have twice the pixel size and cover half the (linear) field of our present CD47-20 CCDs. As a result they have 16 times fewer pixels. They are NIMO parts unlike the AIMO CCDs we currently operate. This, taken together with their smaller parallel clock capacitance, makes their parallel clock speed an order of magnitude faster as well. We anticipate frame rates of ~1 KHz in fast or pipelined occultation modes for ~30 arcsecond square subframes binned 2x2. This is more than sufficient to resolve all of the telescope structural modes and the 300 Hz oscillation of the chopping secondary control system with continuous data. Finally, these CCDs are fast enough to allow early experimentation with a closed-loop image motion compensation system using the analog inputs to the chopping secondary mirror as the actuator in the system.

ACKNOWLEDGEMENTS

HIPO development is supported under USRA Grant 8500-98-003. We thank Bob Leach for the binning and subframing infrastructure provided in the DSP code from ARC. This forms the underpinnings of the HIPO readout software. The idea for fast occultation mode came up many years ago in a conversation with Lloyd Robinson. Paul Jorden provided advice and suggestions for CCD operation on numerous occasions. We also thank Phil Massey and Steve Schechtman for helpful discussions regarding their experiences with ground-based instrument throughput measurements using spectrophotometric standard stars. Finally, we thank Lowell's Director, Bob Millis, for his unwavering support of the project.

REFERENCES

- [1] Elliot, J.L., "Stellar Occultation Studies of the Solar System", *Ann. Rev. Astron. Astrophys.* 17, 445-475 (1979).
- [2] Charbonneau, D., T.M. Brown, A. Burrows, and G. Laughlin, "When Extrasolar Planets Transit Their Parent Stars", in *Protostars and Planets V*, University of Arizona Press, 701-716 (2007).
- [3] Dunham, E.W., J.L. Elliot, T.A. Bida, and B.W. Taylor, "HIPO – A High-speed Imaging Photometer for Occultations", *Proc. SPIE* 5492, 592-603 (2004).
- [4] Dunham, E.W., "The Optical Design of HIPO: A High-speed Imaging Photometer for Occultations", *Proc. SPIE* 4857, 62-72, (2002).
- [5] Taylor, B.W., E.W. Dunham, and J.L. Elliot, "Performance of the Lowell Observatory Instrumentation System", *Proc. SPIE* 5496, 446-454, (2004).
- [6] Leach, R.W., F.L. Beale, and J.E. Eriksen, "New-generation CCD controller requirements and an example: The San Diego State University generation II controller", *Proc. SPIE* 3355, 512-519 (1998).
- [7] Dunham, E.W., R.L. Baron, J.L. Elliot, J.V. Vallergera, J.P. Doty, and G.R. Ricker, "A High Speed, Dual-CCD Imaging Photometer", *Pub. A.S.P.*, 97, 1196-1204 (1985).
- [8] Elliot, J.L., *et al.*, "Image Quality on the Kuiper Airborne Observatory. I. Results of the First Flight Series", *Pub. A.S.P.* 101, 737-764 (1989).