The Slitmask Alignment Tool: robust, efficient, and astronomer-friendly software for aligning multi-object slitmasks

Marc Kassis*[a], Gregory D. Wirth[a], Andrew C. Phillips[b], Charles C. Steidel[c]
[a]W. M. Keck Observatory, 65-1120 Mamalahoa Hwy, Kamuela, CA 96743;
[b]UCO/Lick Observatory, University of California, Santa Cruz, CA; [C]California Institute of Technology, Pasadena, CA

ABSTRACT

Multi-object spectroscopy via custom slitmasks is a key capability on three instruments at the W. M. Keck Observatory. Before observers can acquire spectra they must complete a complex procedure to align each slit with its corresponding science target. We developed the Slitmask Alignment Tool (SAT), to replace a complex, inefficient, and error-prone slitmask alignment process that often resulted in lost sky time for novice and experienced observers alike.

The SAT accomplishes rapid initial mask alignment, prevents field misidentification, accurately predicts alignment box image locations, corrects for flexure-induced image displacement, verifies the instrument and exposure configuration, and accommodates both rectangular and trapezoidal alignment box shapes. The SAT is designed to lead observers through the alignment process and coordinate image acquisition with instrument and telescope moves to improve efficiencies. By simplifying the process to just a few mouse clicks, the SAT enables even novice observers to achieve robust, efficient, and accurate alignment of slitmasks on all three Keck instruments supporting multislit spectroscopy, saving substantial observing time.

Keywords: multislit spectroscopy, slitmask alignment, observing tools, on-sky efficiencies, W. M. Keck Observatory

1. INTRODUCTION

The W. M. Keck Observatory instrument suite features three multi-object spectrographs that can obtain spectra of numerous objects simultaneously. In advance of each observing session, observers using the optical spectrographs DEIMOS[1] and LRIS[2-5] or the infrared spectrograph MOSFIRE[6] design slitmasks that are intended to pass only the light from the science targets of interest, blocking neighboring objects that would otherwise contaminate the data. Each mask is designed with coordinates and a position angle that is specific to the observer's target field, and includes two kinds of elements: *alignment boxes* corresponding to stars of known celestial position which are used to register the mask on the sky, and *slits* which coincide with the science targets.

Slitmasks for Keck's two optical multi-object spectrographs are fabricated on-site using a CNC mill to cut slits and boxes in thin sheets of aluminum. Those for LRIS typically include 5–40 slits, while DEIMOS masks can accommodate up to 150 slits by virtue of the greater available slit length on this instrument. When narrowband filters are employed to limit spectral range, multiple ranks of slits can be placed on a single mask, allowing 1,000 in the case of bright sources. Observers submit slitmask designs to a central server weeks in advance of their observing run to ensure that observatory technicians can mill the masks in time. The observing team can request a maximum of nine to twelve slitmasks to be loaded into the instrument on any given night. In contrast, Keck's new infrared MOSFIRE spectrograph employs a completely different approach: instead of inserting premilled masks into the cryogenic dewar, it employs a fully robotic *configurable slitmask unit (CSU)*[6] to deploy 46 slits in real-time. Although no limit is imposed on the number of MOSFIRE "slitmasks" which can be deployed for the night, in practice the number of slit patterns used per night is comparable to that achieved with LRIS and DEIMOS. During the night, these masks are configured in the optical beam one at a time. Once the observer aligns a given mask to the target field on sky, spectroscopic observations can begin.

The procedural goal of aligning a mask is to center the boxes on their respective alignment stars, at which point the slits will align with the much fainter (and thus harder to see) science targets. The basic process involves two distinct phases: *coarse alignment* and *fine alignment* (see Figure 1). The coarse alignment step is intended to roughly align the mask such that the alignment stars are visible within their respective boxes. Since the telescope pointing inaccuracy at Keck is significantly greater than the size of the typical 4 arcsec slitmask alignment boxes, the alignment stars are generally not

Observatory Operations: Strategies, Processes, and Systems IV, edited by
Alison B. Peck, Robert L. Seaman, Fernando Comeron, Proc. of SPIE Vol. 8448,
844807 · © 2012 SPIE · CCC code: 0277-786X/12/$18 · doi: 10.1117/12.926693

coincident with their boxes once the telescope reaches the target field. Historically, the two optical spectrographs at Keck employed quite different methods to achieve coarse alignment. With LRIS, the practice was to first image the target field twice using the science detector in imaging mode (once with the slitmask in place and once without), then measure the position of a single alignment star and its corresponding box, and finally, to determine the telescope offset required to put the alignment box onto the star. Because the process involved acquiring two CCD images and completing a slitmask move this procedure was relatively slow, averaging 11 minutes to complete. In contrast, DEIMOS observers coarsely aligned masks more rapidly by employing custom software that predicted the precise pixel positions at which certain stars should appear in the DEIMOS offset guider field of view when the mask was properly aligned.

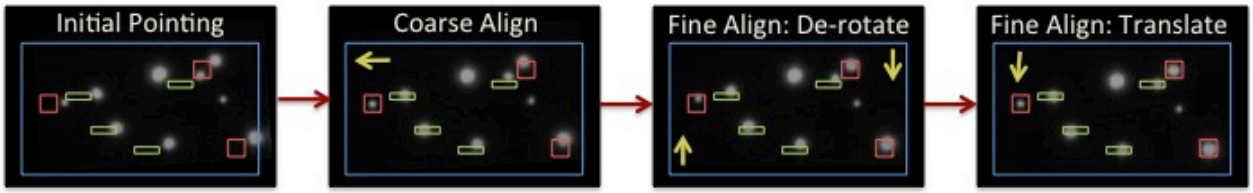

Figure 1. Steps in the mask alignment procedure. (a) After initially pointing the telescope to the target field, the objects are offset from their corresponding slits (yellow rectangles) and box positions (red squares). (b) Coarse alignment involves offsetting the telescope (yellow arrow indicates motion) so that alignment objects enter their respective boxes. (c) Fine alignment involves measuring and applying the instrument rotation angle required put the alignment boxes on the alignment stars. (d) Fine alignment is completed by applying a telescope offset to center all of the stars within their alignment boxes.

With both LRIS and DEIMOS, observers completed the fine alignment process by using *xbox*, an IRAF-based software program written by Drew Phillips for LRIS and later adapted to DEIMOS. Mask images acquired after coarse alignment are fed to *xbox,* which measures the locations of the stars relative to the box centers and thus derives the translational and rotational moves necessary to center the slitmask on the alignment stars. Typically, two to three fine alignment iterations were adequate to align the mask within both rotational and translational tolerances.

The coarse alignment methods and *xbox* program for fine alignment have proven immensely useful for over a decade, but significant room for improvement in the procedure existed. In particular, we desired a method which was easier for observers to master and which could eliminate the numerous procedural errors that commonly cost observers precious observing time. These common problems include (but are not limited to): positioning the telescope to the wrong coordinates or rotation relative to the designed mask; inserting the wrong slitmask; leaving a dispersive optic or the wrong filter in the optical path; misreading and mistyping coordinates measured by eye off images and finder charts; entering incorrect exposure parameters; incorrectly specifying the coordinates for the alignment boxes; misidentifying the alignment stars; and deviating from the prescribed procedure.

These errors and the learning curve produced significant losses of sky time. Using the old methods, most observers could align slitmasks within about 11 minutes, whereas observatory support staff and expert observers could typically complete the alignment process in 5-6 min. To put the cost of this inefficiency in perspective, consider that slitmask-enabled spectrographs are currently used a total of 160 nights over a Keck observing semester. Assuming five mask alignments are completed per night at an average of ten minutes per alignment, the amount of time spent positioning the telescope over the course of a calendar year is 20 nights. A more efficient slitmask alignment procedure that all observers could execute as if they were experts would thus save roughly half that time, potentially making 10 additional nights of observing time available to science – quite a substantial savings.

A hardware upgrade to the LRIS guiding system and the impending arrival of MOSFIRE further motivated the support staff to upgrade the alignment software. First, we needed the slitmask alignment procedure to accommodate the typical infrared observing practices that we would use with MOSFIRE. In addition, we wanted to avoid the time-consuming process of switching the MOSFIRE CSU between full-field imaging and slitmask modes, as would be required if we used the LRIS-style imaging coarse acquisition method on MOSFIRE. Recent upgrades to the guiding hardware on board LRIS have improved the stability of the guider, potentially allowing LRIS observers to employ the much faster guider-based coarse alignment procedure. Thus, for both MOSFIRE and LRIS we wanted to develop DEIMOS-style guider coarse alignment software to dramatically improve efficiencies.

With these motivations, the Keck Observatory support astronomers developed the efficient and observer-friendly Slitmask Alignment Tool (SAT) to replace existing procedures on LRIS and DEIMOS and to enable alignment of MOSFIRE masks. This tool is currently in use on LRIS and MOSFIRE and is being deployed on DEIMOS. Because the SAT will be common to all three instruments, observers will now only need to master one tool to become proficient in mask alignment on all of these spectrometers. In the sections below, we describe key aspects of the SAT interface, the algorithms employed for box and star fitting, and the modifications required to support infrared observing with MOSFIRE. In the last two sections, we characterize the efficiency improvements resulting from SAT deployment at Keck and detail how to adapt the SAT to support multi-object spectrographs on other telescopes.

2. SLITMASK ALIGNMENT WITH THE SAT

The SAT is an application written in the Interactive Data Language (IDL) featuring a user-friendly graphical interface depicted in Figure 2. The software can control instruments and telescopes at Keck by a combination of shell scripts and keyword calls spawned from IDL. The GUI layout was expressly designed to lead an observer through the alignment process in an efficient manner. Clickable buttons are strategically placed to help provide visual clues to the next step in the procedure.

The primary SAT functions are divided into tabs that observers use to toggle back and forth between various modes. When presented in these "bite-sized" screens, the alignment process is clearer and observers find it relatively easy to learn the few steps each tab screen requires to complete the function. The order of the tabs reflects the order of operations, with coarse alignment options listed first and fine alignment listed second.

We have made a deliberate effort to generalize the SAT functions and procedures, which has proven challenging because each instrument is unique. To this end, we strive to record all instrument-specific parameters in arrays that contain information such as the detector geometry, pixel size, field of view, alignment box size, available imaging filters, relative astrometry between the guider and science field, guider characteristics, relevant computer names, etc. Exposure parameters and instrument configurations are stored in these arrays and default values are used on startup.

2.1 Coarse alignment

The goal of the coarse alignment phase is to align the mask just well enough that all of the alignment stars are visible in their corresponding alignment boxes on the slitmask. Since the Keck guiders read out far faster than the science detectors, this step can be completed most rapidly if the observer can determine the required offsets from a guider image. If the astrometric position and orientation of the guider field of view relative to the slitmask center are well determined, then the approximate RA/Dec offsets required to coarse-align the mask are easily determined by comparing a guider image to a predicted image that would be obtained if the telescope were pointed to the correct position. The Guider Coarse Align tab on the SAT enables this comparison.

When launched, the SAT defaults to the coarse alignment screen on the assumption that this is the first step an observer will complete. The upper left of the coarse screen holds functions and an image display that are designed to predict the guider field (see Figure 2). Observers have the option of loading a target list that includes the desired rotator position, and may use the same target list that is employed by the telescope and guiding software. Loading the target list is completed once during the night and populates the target drop list. For each mask alignment, observers follow the procedure below to coarse-align a mask using a four-click process:

1. *Select Target* using the target drop list. When the target is selected, the GUI:
    - Queries a DSS image server to obtain an image of the guider field corresponding to the science target.
    - Displays the DSS image in the upper left of the coarse align screen.
    - Overlays red boxes on the six brightest objects in the field obtained from the USNO catalog.
    - Displays a table of objects names, celestial coordinates, and predicted guider pixel locations.
    - Erases any existing offsets from the screen as a visual clue to the observer.
    - Disable *Move Telescope* button to protect the observer from triggering an inappropriate move.

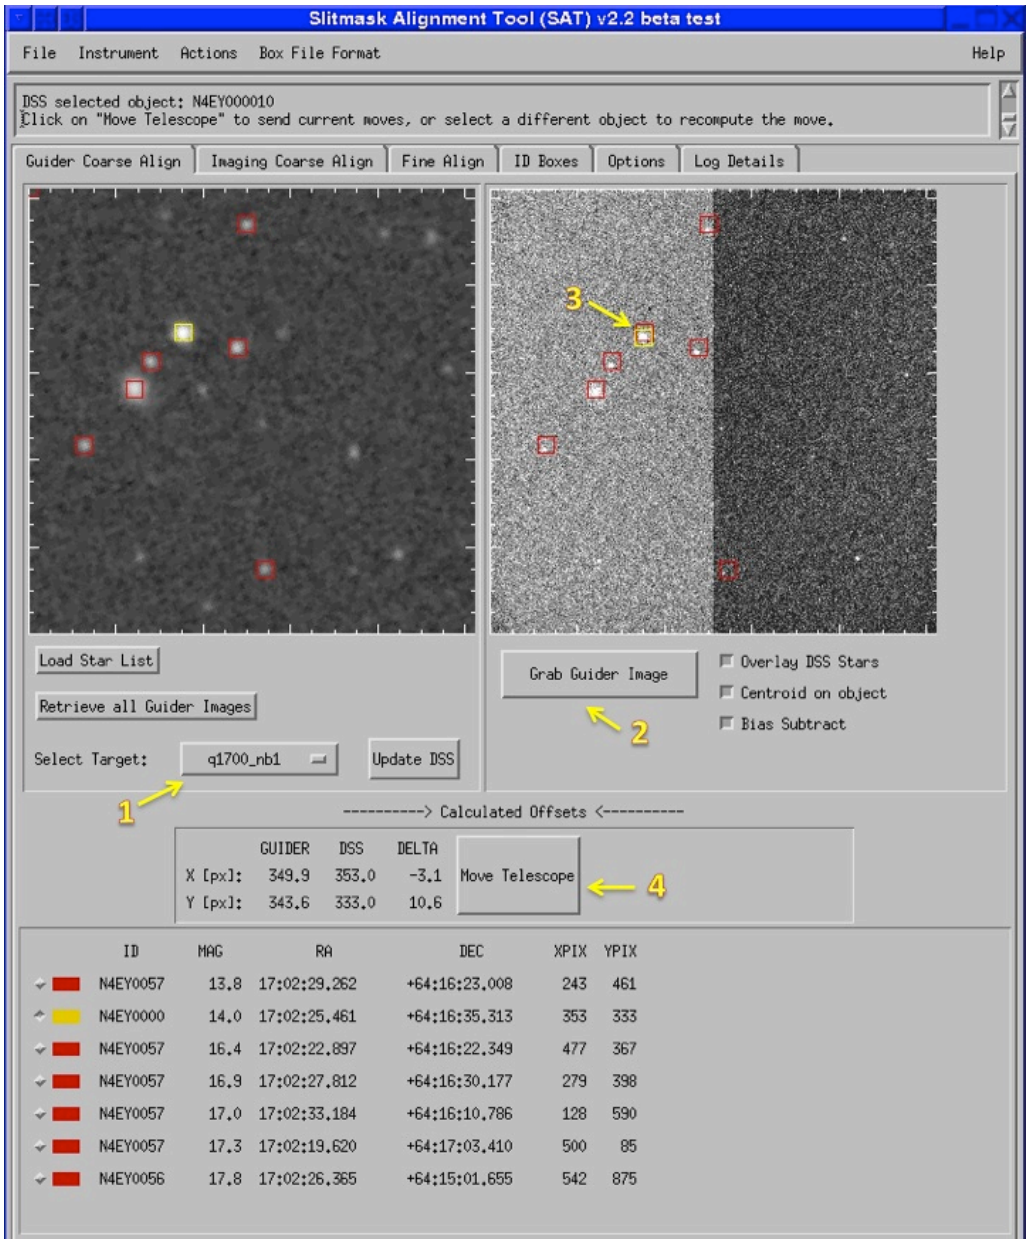

Figure 2. SAT coarse alignment screen showing a predicted DSS (left) and current guider (right) image corresponding to the selected target. The arrows labeled 1-4 indicate the four clicks necessary to coarse align a slitmask: (1) select the target; (2) save and display a guider image; (3) identify the alignment object; (4) send a telescope move. See text for additional details regarding the actions taken at each mouse click.

2. Click *Grab Guider Image* on the right hand side of the screen to trigger a guider image acquisition, which:

- Saves and display a guider image in the upper right display.
- Overlays red boxes at the predicted positions of the guide stars.
- Finds the asterism of the seven alignment objects in the guider image as described in section 3.1.
- Overlays white boxes on measured object positions that match the seven-object asterism.
- Disables the *Move Telescope* button to protect the observer.

3. Click on an alignment object in the displayed guider image and the SAT:
    - Measures the position of the object.
    - Overlays a yellow box on the star to verify the selection and centering.
    - Locates and marks the corresponding star in the DSS image.
    - Toggles the selection indicator from red to yellow in the table for the nearest star.
    - Calculates the offsets and display them in the offsets table below the image.
    - Activates the *Move Telescope* button to lead observers to the next step.
4. Click *Move Telescope* and the SAT:
    - Offsets the telescope using the computed offsets.
    - Reacquires and displays a guider image as a double check that the telescope move was correct.

At this point, the mask is coarsely aligned with objects in the alignment boxes (Figure 1: Coarse Align), and the observer may proceed to the fine alignment phase. Before we leave the coarse alignment phase, there are a couple of useful SAT features worth mentioning. The first is using the *Retrieve All Guider Images* option to retrieve all the required DSS images for their target fields in the observer's starlist and store them on disk. Retrieving DSS images in advance will save time on sky, but is not done automatically for observers because target lists may be in flux during the first part of the afternoon or night when the SAT is started. Observers also have the option of using the current telescope position as the target. This is useful for calibration stars or targets of opportunity that may not be in the observer's target list at the time of observation. Last, if the yellow-highlighted nearest star in the DSS image is not the correct object to use for coarse alignment, a different alignment object may be selected either by clicking a different object in the DSS image or selecting a different object in the alignment star table.

2.2 Fine Alignment

Once coarse alignment is complete, the remaining task is to refine the telescope pointing and the instrument rotation angle in order to center the slitmask alignment boxes optimally on their corresponding alignment stars. This requires measuring accurately the offset of each alignment star from its corresponding box center and determining the offset and rotation that will optimize the positioning. To do so, observers navigate to the fine alignment tab of SAT to optimize functions that will refine the positions of the stars within their respective alignment boxes and achieve optimal centering of the slitmask. When observers first navigate to the tab, the graphs and offset table are blank. On the left side of the fine alignment tab, box plots are presented for each alignment object as seen in Figure 3. There are two plots that summarize the fit along the x and y-axis of the image, respectively. At the upper right, a residual plot marks the locations of the boxes and vectors indicating the magnitude of motion needed to correct each star to the desired location. The bottom right of the tab pane presents the recommended moves and provides several options to the observer.

Fine alignment is basically a two-step process. The first is to acquire and analyze the images, while the second is to send the recommended moves. Below we present the steps to fine aligning a mask.

1. Click *Start Fine Alignment*
    - Reconfigures the instrument if necessary to acquire mask images.
    - Warns observes if the slitmask is not in beam.
    - Triggers an image acquisition using the desired instrument and exposure settings. The current exposure is used if the acquisition was started without the SAT via the standard instrument exposure controls.
    - Waits for the exposure to complete, reads the resulting image, then computes box and object centers and the translational and rotational offsets needed to align the mask.
    - Updates the box plots residual map and the offset table with the information necessary to assess the fit.
    - Recommends rotational and translational moves by setting the default move status to "Y" or "N," depending on whether the move exceeds the significance threshold.

2. Click either *Send and Retake Image* if the observer wants another alignment iteration, or *Send Moves Only* if the observer feels the fit has converged and alignment will be satisfactory once the suggested moves are sent.
    - Applies selected translational and rotational offsets
    - If the retake option is used
        - Acquires image and return to step 1 above
        - Erases the plots and calculated offsets as a visual clue that the analysis is incomplete.
3. (optional) Repeat step 2 until fit has converged, and then proceed with science observations. Convergence is typically achieved within two iterations.

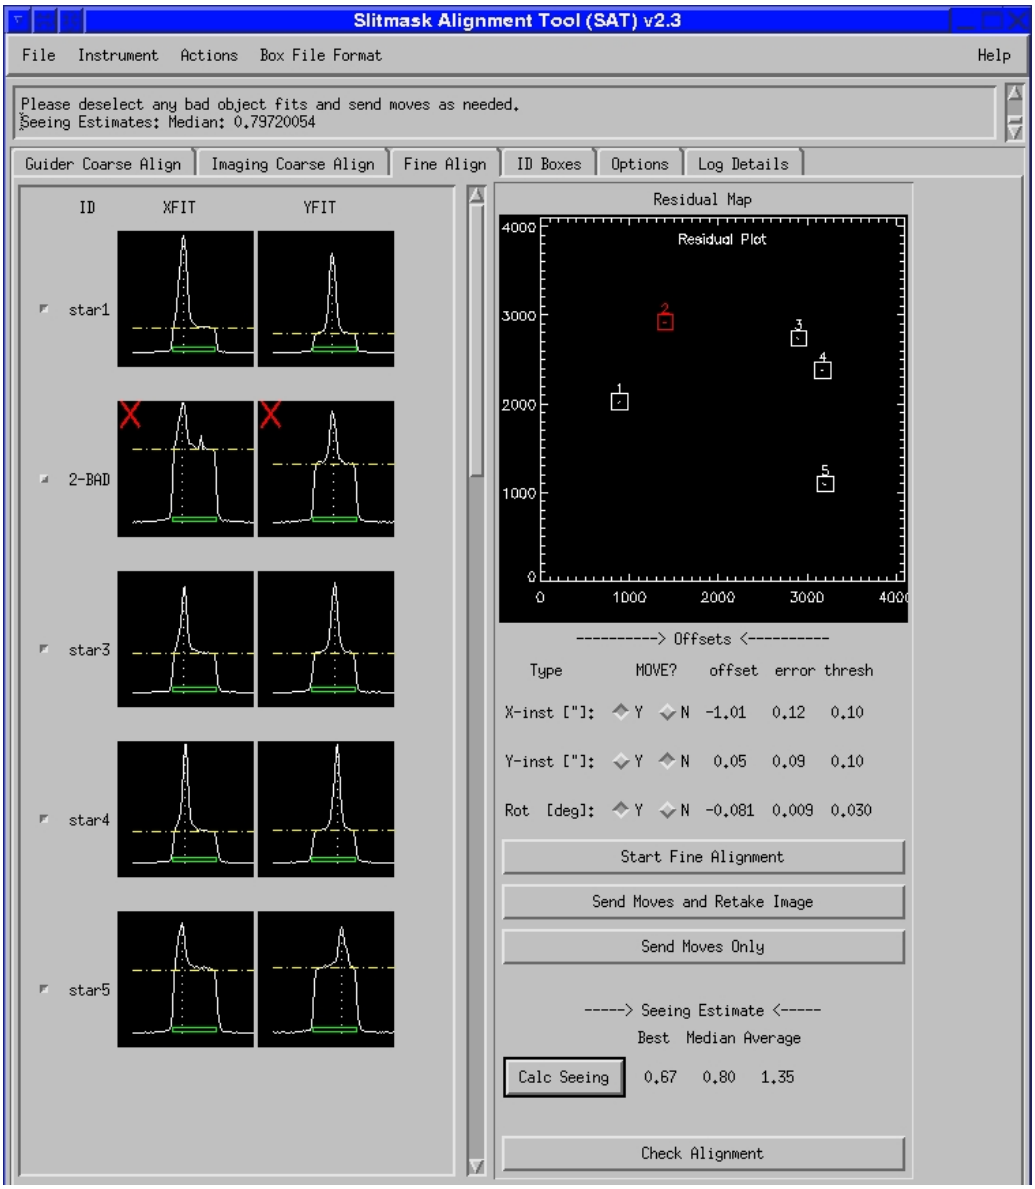

Figure 3. SAT fine alignment tab pane displaying the output for an LRIS mask image acquired when the observer clicked *Start Fine Alignment*. The box plots on the left of the screen display the normalized box profile (white line), the measured position of a box (green rectangle at the bottom of each plot), the measured center of the object (vertical dotted line), and the minimum sky emission in the box (yellow

dash-dot line). As an example, object 2 is flagged as "bad" by the observer to indicate that its centroid is suspect. The residual plot in the upper right shows the locations of the boxes in the field of view and vectors indicating the motion each box would need to move for optimal centering. The area below the residual plot displays recommended telescope moves.

After each image acquisition and analysis in the steps above, the observer must assess the quality of the fit and make a decision concerning the convergence of the fit. This may require one or more iterations as the solution converges, and all of the tools necessary to assess convergence are at the observer's fingertips.

Observers should review the box plots and the residual plots to check for outliers. If a fit for a particular object is suspect, the observer may choose to ignore the object by clicking the toggle button located to the left of the box plots. When a box is toggled from the fit, the SAT recalculates the fit excluding the object. It also marks the box plots as bad in two ways: the label "bad" appears next to the toggle box and red Xs appear in the upper left of each box plot. In addition, the box is marked red in the residual plot (as shown in Figure 3). The residual plot labels each box with the object number displayed to the left to aid object identification. More than one object may be ignored in a fit. If only one box is considered "good," only translational offsets can be calculated.

The SAT displays offsets and recommended moves below the residual plot in units of arcseconds for translational moves and degrees for rotation. The uncertainties in the calculated offsets are presented. The last column in the offset table provides the threshold for making a move. For LRIS and MOSFIRE, the thresholds for making moves are 0.1 arcseconds in translation and 0.03° of rotation. Because the DEIMOS FOV is significantly larger, the tolerance on rotation is tighter (0.01°) to ensure that objects at the edges of the mask are in the slit.

If a move exceeds the tolerance, the *Y* button in the offset table is preselected to indicate a move is recommended. Observers may decide to not send moves by clicking the *N* option. Observers may elect not to send moves when the error in the measurement is large or the image quality is poor or unstable. Under typical observing conditions, observers opt to send the recommended moves.

Below the offset table is information on the seeing calculated using the objects in the alignment boxes. Average, median, and best seeing estimates are provided in the table. Occasionally, extended objects (e.g., galaxies) are used as alignment objects, so the median and best seeing estimates are the preferred estimators of the characteristic seeing. This quick look at the data is useful to observers who are assessing in real time how long to observe their field. Observers who notice large seeing values may decide to observe a field longer than initially planned in order to achieve the desired signal-to-noise.

3. CALCULATING OFFSETS

At the heart of calculating the translational and rotational offsets to align the slitmask is a fast and effective box-fitting algorithm originally written by Drew Phillips as an IRAF routine called *xbox* and translated into IDL. The SAT version of the routine requires as inputs the image of the mask and the predicted locations for the box positions on the mask. We have augmented the original fitting process with a flexure compensation algorithm that can account for instrument flexure or deviations in the optical path due to different optical elements being in beam. In the next few sections we describe the fitting process, improvements made to the original routine, and the unique modifications needed for observing with MOSFIRE.

3.1 Automated flexure correction algorithm

A key barrier to fully automating the slitmask alignment process is image motion that changes the positions of the alignment boxes. Both LRIS and DEIMOS are subject to gravity-induced flexure, which can produce image motions greater than the typical box width. Furthermore, the image position on LRIS depends on which dichroic is used; as a result, if the observer uses a different dichroic when measuring the box positions in the afternoon and aligning the mask on-sky, the boxes will not appear at the predicted location. These two image motions can be large enough to prevent the box-measuring algorithm from functioning reliably at all instrument orientations unless the software corrects the predicted box positions. To serve this need, we developed a simple and novel technique for the correction of image motion, which is both fast and reliable.

Given the predicted positions of the alignment boxes and a maximum expected amount of flexure, the algorithm operates as follows. First, we extract a square region centered on the predicted location of each box, sized to include the maximum amount of image motion expected. Within that region, we set all pixels above a given threshold value (selected to divide illuminated and unilluminated pixels) to 1, and those below the threshold to zero, thereby forming a "submask" which identifies illuminated pixels that may belong to a box. We then "stack" these identically-sizes submasks to form a "super-mask" and locate pixels that have a value of *N*, where *N* is the number of boxes expected. The key is that the *actual* position of each box is offset from the *predicted* box position by the same amount on each submask; hence, there will generally be one and only one contiguous region of pixels with value *N* on the resulting supermask. A simple marginal centroiding algorithm applied to the supermask returns the (x,y) center of the common region; measuring the offset from this (x,y) position to the center of the extraction box then yields the offset of the boxes relative to their expected positions. Adding this offset to the predicted box positions will shift them to coincide with the actual positions.

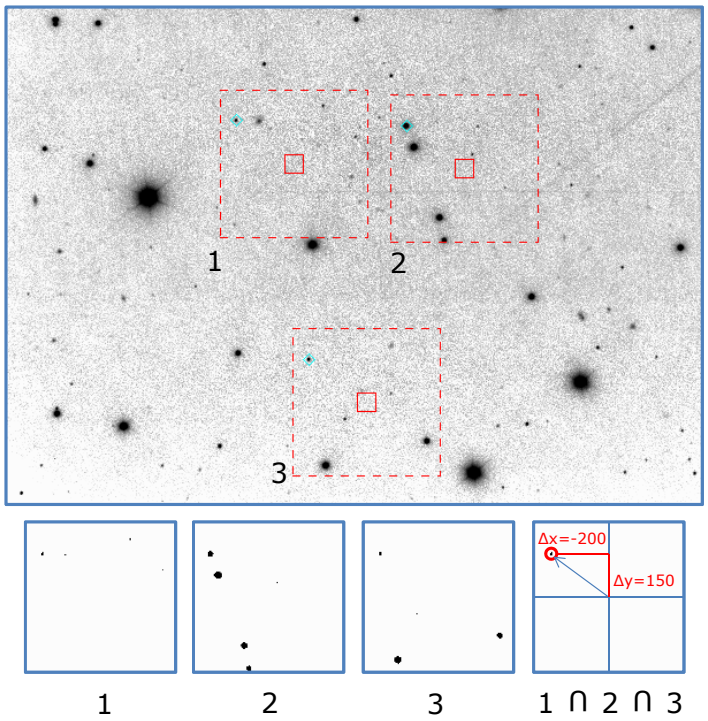

Figure 4. Example showing the operation of the flexure correction algorithm applied to determine of the telescope offset. Top panel: view of a sky region showing the locations of three alignment stars (black diamonds) which are offset from their corresponding slitmask alignment boxes (solid red squares) due to poor telescope pointing. The dotted red squares indicate the size of the 256x256 search region around each box. Lower panels display the corresponding illuminated pixel masks for each of the three search regions, where pixels significantly above the local sky background are shown as black (1), those below as white (0). The rightmost lower panel shows the intersection of the three search regions. Only one set of pixels is illuminated in all three search regions, allowing the algorithm to determine that the offset between the box and star centers is -200 pix in X and +150 pix in Y.

The beauty of this technique is that the very same algorithm can be employed to automate two other aspects of the mask alignment process. First, it can be used in the guider coarse alignment method to determine the offset between the predicted and actual positions of the guide stars. Second, in the imaging coarse alignment method it can identify the offset between the alignment stars and their corresponding boxes. Its ability to function well even in relatively crowded fields makes it a particularly useful technique in this case, since it can identify the correct alignment star pattern much more effectively than the human operator.

3.2 Box and star centering algorithm

For each refined box position, the routine identifies a sub image that now includes the image of the box, star, and the surrounding area of the mask that is not exposed to light. A typical sub image has a size that is equivalent to twice the box width in both directions. Example box images for LRIS and MOSFIRE are shown in Figure 5.

To determine the box and star centers, we first derive the box profiles by summing the image rows and columns in the vicinity of the box. Derivatives of the box profiles are next calculated and compared against model derivatives for a box and star. Due to finite sampling, the alignment box edges are "soft," so that the box profile has sloping sides instead of an ideal "top hat" profile. The box model derivatives try to match the emission at the box edges and assume where there is no emission or sky emission that the derivative is zero (see figure 5). The box center is the position at which the model and measured derivatives have the highest correlation.

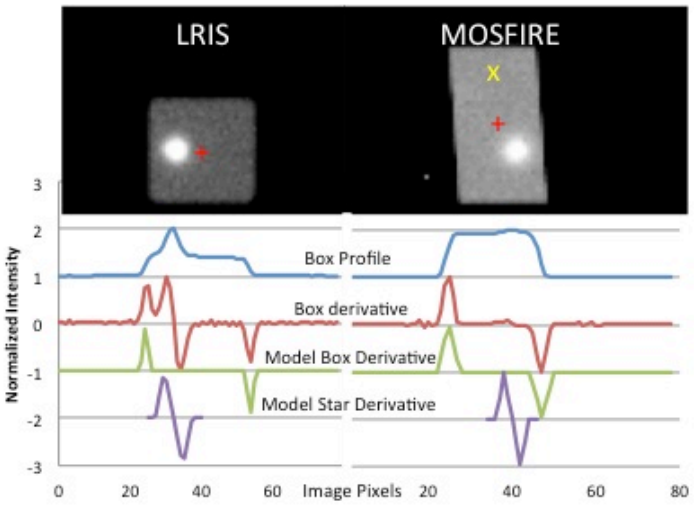

Figure 5. Images of typical LRIS (left) and MOSFIRE (right) alignment boxes with stars offset from the desired position, indicated as the red "+" in the left-hand image and as the yellow "X" on the right-hand side. Below each image are plots of several quantities important in the calculation of the box and star centers: the box profile along the x-axis, the derivative of the box profile, and the model derivatives for a box and star profile. The plots were normalized to the peak intensity and then an offset was added.

To fine-tune the measured center of a box, the routine fits each box twice. On the initial pass, the routine uses the flexure-corrected box position to extract a subimage and measure a preliminary center for the box. On the second pass, the algorithm extracts a new sub-image centered on this preliminary center. This second pass ensures that the box center is measured using a sub-image that completely samples the emission inside the box.

To determine the alignment star position, the routine fits a derivative of a box of no width where the edges of the box are matched to the typical seeing. The derivative of the object model is shown as the bottom line of figure 5. Again, two passes are made to fit the star: the first to identify the star, the second to fine-tune the position. The end result of this procedure is an accurate measurement of the offset of each star from the center of its box.

3.3 Calculating the offset

The next step is to project the box and object centers onto a uniform coordinate system. The detectors for both LRIS and DEIMOS consist of a mosaic of CCDs with different layout geometries. Parameters in the instrument-specific state files contain information about the CCD geometry and include information such as the CCD size, mosaic layout (linear or two dimensional), and the gap sizes between the CCDs. An understanding of the geometry is not critical for translational moves, but for rotation the offset will not be accurate without knowledge of the true lever arm between the center of rotation and measured position of a star and box pair.

For each instrument a center of rotation is also defined. The center of rotation is typically defined to be the mask field center, and this position may differ from the center of the field of view, as is the case for all three multi-object spectrographs at Keck. The rotational center is then defined as the origin for calculating offsets via a least-squares fitting algorithm. The results are then displayed to the SAT fine alignment screen as depicted in Figure 3.

3.4 MOSFIRE modifications

Until April 2012, our experience with aligning slitmasks was primarily at optical wavelengths with LRIS and DEIMOS. Infrared mask alignments with MOSFIRE required modifications to the fitting process to accommodate for infrared mask alignments in general, as well as more instrument specific changes. Below we describe these changes.

3.4.1 Box fitting modifications

In the infrared, sky emission often dominates the signal from the alignment object, thus it becomes necessary to subtract a "sky image" (acquired at a small spatial offset from the target position) to detect an object in a box. As an example, the stellar contribution to the total flux is clearly visible for LRIS in figure 5, but for a MOSFIRE image, the sky emission dominates. We tried two approaches to subtracting a sky. The first was to average the lower quartile of the emission found within a box and subtract that value from the box sub-image. This method failed for relatively faint objects because the noise from the pixel-to-pixel variations confused the centering code. We eventually settled on a procedure that first acquired a "sky" image that was offset from the mask target position by a few arcseconds. This is automatically accomplished when observers click the *Start Fine Alignment* button to begin mask alignments. If the current sky offset parameters are less than the size of a box, then the SAT will default to offsets large enough to avoid imaging the object in the box at a sky position. The sky image is acquired first so the measured object positions of the stars correspond to the current telescope position. For determining box centers, a single image is used, but when measuring object centers, SAT subtracts the sky image from the on-target image to remove the background. When additional images are needed to converge to a fit, the SAT will only acquire the on-target image and subtract the same sky acquired in the initial exposure set.

As indicated in Figure 5, MOSFIRE's boxes differ from those in LRIS and DEIMOS in two ways: (1) the boxes are longer in one dimension because the bar widths are a constant 7.0 arcsec in width (vs. 4.0 arcsec on LRIS and DEIMOS); (2) the edges of the slits are slanted relative to the detector rows and columns to improve background sampling and this yield superior sky subtraction. Fortunately, minor adjustments to the box fitting parameters can accommodate both of these differences within the existing algorithm developed for DEIMOS and LRIS. In particular, the parameter describing the softness of the box edge is increased slightly to account for the more trapezoidal box profile shape (see Figure 5). Note that the widths of the peaks in the observed and model derivatives for the boxes are wider for MOSFIRE than for LRIS.

An additional and important distinction is that MOSFIRE's alignment objects are generally offset along the length of the box rather than lying at the center because the bars on MOSFIRE have a fixed width and y-position. To compensate for this offset, box centers are first calculated, and then the offset is applied to the box centers so that the SAT knows where the alignment objects should be centered. Figure 5 shows the position from center where a star should be moved for the mask to be aligned, and this offset is reflected in the box plot example in Figure 6.

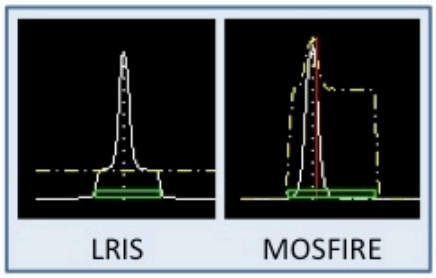

Figure 6. Box plots summarizing the modifications required for MOSFIRE. For MOSFIRE, SAT displays the box profile (yellow dash-dot line) and the object profile after sky subtraction (white line) while LRIS only displays the box profile (white line). A vertical red line indicates the position where an object should be centered, while for LRIS the object center should be the same as the box center, marked by a green rectangle at the bottom of the plot.

3.4.2 Long slit alignments

With the two optical spectrographs, single object acquisition is accomplished in a straightforward manner using an optical guider which can directly image the long slit. Thus, for single object acquisitions, we simply identify the object in the guider FOV, measure the centroid of that object, and then offset the telescope to move the slit onto the target. Since MOSFIRE lacks a slit-viewing guider, single object acquisition is accomplished using the science camera and the

SAT. This works just like a typical fine alignment on a multi-object mask, but in the analysis, only the translational offsets are calculated. This is an advantage MOSFIRE has over LRIS and DEIMOS, in that long slits may be positioned any where in the field a normal slit may be configured. With LRIS and DEIMOS, alignments are limited to the long slit position as seen by the guider.

4. OTHER TOOLS

In addition to the two primary views for coarse and fine aligning a slitmask, there are other useful functions built into the SAT that benefit both observers and support staff. These useful tools include box identification, quick alignment checks, and instrument options.

4.1 Align Check

In order to obtain spectra of sufficient S/N on faint targets, observers may integrate on a given slitmask field for several hours, during which time a spectrograph mounted at the Cassegrain or Nasmyth telescope focus may rotate by many degrees in order to maintain the correct position angle of the mask on the sky. This rotation tends to produce flexure between the guider and the slitmask, which can ruin the mask alignment. For such programs it is necessary to interrupt data-taking periodically in order to acquire an alignment image to confirm that the mask remains properly aligned with the science targets. A *Check Alignment* option is provided to the user on the bottom of the alignment screen (see Figure 3). When *Check Alignment* is clicked, the SAT launches a routine that performs a quick check of the mask alignment under the assumption that the shift will be small and that the alignment stars will still be in their boxes.

The routine waits for the current image to finish exposing, and then while the spectroscopic image is reading out, reconfigures the instrument for imaging and sets the appropriate exposure parameters. Once the instrument is reconfigured, the script acquires a single alignment image and then immediately reconfigures the instrument for spectroscopy while the image is being read out and written to disk. The fine alignment analysis is performed on the resulting image of the mask, and the observer is given the option of performing the derived translation and rotation if desired. The observer then initiates a new science image using the instrument control software.

4.2 Box Identification

In order to perform the fine alignment analysis, the SAT needs to know approximately where to locate the alignment boxes in the images. This information is provided in different ways for each instrument. LRIS observers must create x,y box coordinate lists for each slitmask. The coordinate lists are then used as the initial guess for the position of the boxes during the fine alignment process. Images of slitmasks are typically acquired in the afternoon and then displayed in the SAT *ID Box* image screen. Observers then position the cursor near an object and either mark a box or delete a box nearest to the cursor. A centering algorithm is used to determine the rough center of the box if the option to centroid is selected (this is the default setting). Every time a box is marked (“m”) or deleted (“d”) a file containing the x,y image coordinates for the box is saved in the data directory; the name of the slitmask is included in the filename so that the box file is easily associated with a specific slitmask. The SAT keeps track of the box files so observers do not make mistakes by associating slitmask images with the incorrect box coordinate lists.

For both MOSFIRE and DEIMOS, alignment star and box information is stored in a FITS header extension that is included with every image. The box information is parsed by instrument-specific code to predict where boxes will appear in the image. Using the *ID Boxes* screen, observers can overlay the predicted positions of the boxes on a slitmask; this ability to compare the predicted and actual box locations has proven to be a valuable tool when troubleshooting images for all instruments when objects and/or boxes are not correctly identified.

In the case where FITS extensions are not properly being added to the image headers for MOSFIRE and DEIMOS, the observer has the option of using a box file instead of the mask information in the FITS headers. This option is a valuable workaround and troubleshooting tool.

4.3 Instrument-specific options

When acquiring mask images, observers may accept the default instrument configuration and exposure settings or select different values. The instrument filter complement and exposure options may be changed from their default state using available dropdown menus. Exposure options such as exposure time, co-additions, sampling mode, detector windowing,

and detector binning may also be changed. There is also an option to toggle on and off acquisition of sky images. This list of user-selectable options provides observers with flexibility to adapt to different observing conditions or alignment star brightness.

The LRIS and DEIMOS default exposure time is 20 s, which is adequate for most alignment stars. During commissioning, typical MOSFIRE alignment images were single 10 s exposures using a single CDS readout, and we have adopted these exposure parameters as the defaults. Because the readout time is minimal, exposures complete more quickly on MOSFIRE than on the optical spectrographs. Although the sky subtraction option has only been used in the infrared (where it is set by default) so far, we plan to test sky subtraction on the optical instruments in an attempt to improve the analysis when relatively faint alignment stars are used.

5. EFFICIENCY IMPROVEMENTS

During nighttime operations, the SAT is an efficient tool for aligning slitmasks. The learning curve for new observers is gentle, and most are comfortable using the software after aligning their first mask. The predicted amount of time required to align a mask under typical conditions is presented in Table 1. Fine alignment assumes that two mask images are analyzed and two telescope moves executed before the fit converges. Ten seconds is budgeted for the observer to assess the fit at each step.

Following mask alignment, observers must reconfigure the instrument for their science observations. For LRIS and DEIMOS, this switch from imaging mode to spectral mode typically involves simply installing a grism or grating. For MOSFIRE, the configuration time assumes that the grating will be moved and the slitmask will be configured to move from the alignment to the spectroscopic setup. MOSFIRE observers wanting to save 90 s may opt not to reconfigure the slitmask, at the cost of fewer slits on science targets. Most observers will likely redeploy the boxes to slits for science targets, and therefore, 120 s is allotted for MOSFIRE reconfigurations. Table 1 shows that the expected amount of time between the start of mask alignment and the first science observation is under five minutes for all three spectrographs.

Table 1. Execution time

Alignment process steps	**Execution times (sec)**		
	LRIS	**DEIMOS**	**MOSFIRE**
Guider coarse alignment	30	30	30
Fine alignment (2 images)	160	145	84
Configure instrument for science	60	60	120
Total execution time	**250**	**235**	**234**

Besides execution time, the other important efficiency improvement the SAT makes is eliminating procedural errors that would result in significant lost time. As an example, mistyping alignment star coordinates during coarse alignment would result in several minutes lost time, as the observers diagnose why stars are not visible in their boxes. With LRIS and its old method of coarse-aligning the mask by imaging the field, five to ten minutes would be lost recovering. These types of procedural errors are common to all three instruments, and the SAT avoids many by doing the following:

1. Eliminating the need for observer finder charts.
2. Checking the telescope position and rotator angle against those of the desired mask and warning the observer of inconsistencies.
3. Calculating coarse alignment moves eliminating mistyped and misread coordinates.
4. Warning the observer if a mask is not in beam.
5. Automatically configuring the instrument for imaging a mask.
6. Reducing the learning curve for novice observers.

7. Resetting the exposure parameters, avoiding exposures parameters last set for spectroscopic observations.
8. Ensuring the correct box coordinates are used as input.
9. Integrating the acquisition and analysis processes to prevent the observer from analyzing the wrong image.

Although we do not specifically track the amount of time spent aligning masks, we can estimate the time required to align masks by measuring the duration between when the telescope stops slewing to a new target and the initiation of the first multi-object science exposure; ideally, this should be equivalent to the execution times in Table 1. Figure 7 shows how mask alignment times have evolved over the last decade of Keck operations and indicates clear improvements to the alignment process. In August 2002, observers started using the recently-commissioned LRIS blue side camera to align masks, which saved time by avoiding lengthy red-side instrument configurations. The migration from the red to blue camera resulted in a savings of several minutes per alignment on average. Before the introduction of the SAT, only one year (2008) had a median configuration time under 10 minutes, and that was an anomaly. In 2008, several teams of observers repeated the same set of masks on consecutive nights; instead of imaging the field to coarse align the mask, the observers sped up alignment by putting the guide star at the same location on the guider as it had been the night before. These observers used the guider to quickly coarse align the masks similar to the method now used with the SAT.

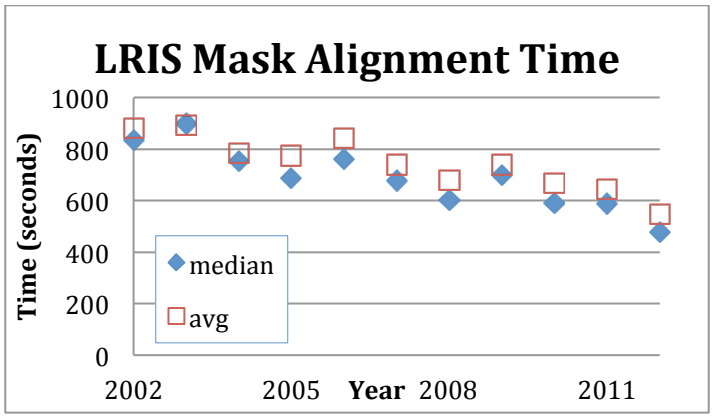

Figure 7. Median and average times required to align LRIS slitmasks. The LRIS red camera was used to align masks in 2001-2002. From 2003-2009, the LRIS blue camera was used for all steps in the alignment process. The SAT entered operation in 2010, with coarse alignment via guider imaging starting in 2012.

The next jump in efficiency occurs in 2010 with the introduction of the SAT. The new LRIS guider hardware was installed in June 2010, and instrument flexure hampered our ability to predict the guider field until the middle of 2011. The year 2012 marks the first full year of coarse aligning with the guider, and the efficiency improvements are noticeable. Comparing the three SAT years to the three previous years (2007-2009), the improvement in efficiency is 2 min per mask. If we compare just the 2012 alignment times to the three years prior to SAT deployment, the time saved is 3 min. Based on the typical number of annual mask alignments, we estimate that SAT is saving us the equivalent of 2.5 nights on-sky, for LRIS alone.

6. PLANNED UPGRADES

Although we have significantly improved the alignment process, room for further improvements remain. In particular, we recognize that there are options to reduce the image size for our optical spectrographs, saving both readout and computational time. For example, alignment images acquired with DEIMOS are automatically windowed to include only the small area of the detector that is illuminated in direct imaging mode. This reduces the readout time to an average of 22 s for a savings of 37 s per image. Although not yet fully deployed on LRIS, we are adapting this technique for LRIS.

We are also considering employing binned data, which will save 17 s and 6 s in readout time for LRIS and DEIMOS, respectively, plus additional computational time. MOSFIRE data cannot acquire binned images. Recently, we acquired unbinned and binned 2×2 image sets, and determined the difference in offset for translation and rotation. For a set of 18 image pairs, the average offset is 0.001±0.106 arcsec in x translation, -0.003±0.085 arcsec in y translation, and 0.001±0.008 deg in rotation. Although, the average agreement is nearly perfect, the translational uncertainties are relatively large compared to the desired accuracy of 0.1 arcsec. We suspect that the model derivatives are not well

matched for binned data, and we are currently fine-tuning the models to improve the fits. We are very hopeful that binning and windowing will become the default option for the optical instruments later in 2012.

For LRIS and DEIMOS we would like to fabricate new long slit masks that contain, in addition to several long slits of varying widths, an alignment box for acquiring targets too faint to be detected in guider images. This would provide observers who want to observe relatively faint single objects with the option of using the SAT to align their objects on the slit instead of trusting a blind offset from a nearby brighter star.

Lastly, we note that the typical mask alignment time has not yet achieved the five minute ideal time predicted in Table 1. Much of the dead time in the alignment process involves waiting for the observer to proceed to the next step. It is our goal to more fully automate the process, with an ultimate goal of achieving "one-click" mask alignments that would accomplish the entire process: acquiring the guider image, measuring and sending the coarse alignment moves, reconfiguring the spectrograph for direct imaging, acquiring the fine alignment images, measuring and sending the translation and rotation moves, iterating until converged within tolerance, and reconfiguring for spectroscopy. We feel that this is an attainable goal for most slitmask observations.

7. ADAPTING THE SAT FOR YOUR MULTI-OBJECT SPECTROGRAPH

The SAT was designed to enhance the efficiency and reliability of slitmask alignment on all of Keck's multi-object spectrographs, and is well on its way to achieving that goal. Since it has already been adapted to work on two telescopes and three different instruments, inquiring minds will ask whether the SAT could be adapted to work on other instruments and other telescopes. We believe that the answer is "yes," with a modest investment of effort. Below, we outline what would be involved in adapting SAT to function with your multi-object spectrograph.

To function with a given instrument and telescope, the SAT requires both a set of instrument parameters (for use in predicting guider fields, configuring the instrument for acquisition images, and calculating offsets) and a set of commands callable from a command shell (for instrument and telescope control). Required parameters for coarse alignment include the guider detector characteristics (pixel scale, detector size, flat fields) and the relative astrometry between the guider and the mask field center (vector to the guider field center, relative guider position angle offset). Required parameters relating to fine alignment include science detector geometry and size, pixel scale, alignment box size, sub-image sizes, typical box edge blur, center of rotation in detector coordinates, mask curvature, and desired tolerances on offset and rotation of the slitmask. Image acquisition, instrument configuration, and exposure control rely on the existence of external commands, specific to a given instrument, which SAT can execute to perform these operations. External scripts used at Keck trigger guider and science camera exposures, move the telescope and rotator, save and restore instrument configurations, and move mechanisms like filter wheels in order to be properly configured for acquiring images of the mask. Whether these functions can be integrated with SAT will depend on whether a given instrument permits control through commands callable from a standard shell.

In addition to the changes mentioned above, certain IDL routines will need to be adapted to work with a new instrument. Each observatory will likely have different standards for formatting target lists, and thus modifications to the target list reader will be required to obtain the coordinates and position angle for each slitmask. If mask data are stored in the FITS images as extensions, the box file interpreter will need to be updated to read the specific format for a given instrument. Finally, a routine must be written to read the images output by the instrument, perform any required processing such as bias correction and assembly of mosaic FITS data, and reformat the image data as needed into a data array suitable for analysis. Given the key role that multi-object spectroscopy plays, and will continue to play, on the world's leading telescopes, the effort required to adapt SAT to function on the world's major telescopes would likely pay very substantial dividends over the lifetime of the observatory. Interested parties are invited to contact the authors for advice on adopting the SAT to their instruments.

8. SUMMARY

Feedback from observers who have used the SAT at Keck indicate that it was one of the highlights of their observing experience. The SAT simplifies the slitmask alignment process by improving automation, orchestrating instrument configurations with image acquisitions, correcting for small displacements of boxes in the images, using the guider

images to coarse align masks, and ensuring that both the instrument and telescope are ready to image the mask. With these improvements, the SAT avoids many of the common procedural errors that resulted in significant time losses, and this savings has endeared it to our observing community.

Using the SAT, observers accomplish the two primary steps for mask alignment for all three multi-object spectrographs at the W. M. Keck Observatory: DEIMOS, LRIS, and MOSFIRE. The user interface is designed to lead observers through the process with text and buttons becoming active at appropriate times to indicate the next step. Optical and infrared observing practices are folded into the SAT, so that the process is the same across instruments, reducing the amount of cross-training required for observers who may use more than one of the multi-object spectrographs to meet their scientific goals.

Since first implemented in 2010, the SAT has reduced LRIS slitmask alignment times by 30%. This translates directly into recovered observing time with a potential savings of six nights over a year's worth of observing sessions for the three instruments. Other observatories may benefit from these savings if their process is prone to procedural errors, and the next generation of telescopes should consider how best to align masks when developing new instruments for those facilities. With modest modification to the IDL based software and investments in engineering time for calibrations, the SAT could be adapted for your multi-object spectrograph.

REFERENCES

[1] Faber, Sandra M.; Phillips, Andrew C.; Kibrick, Robert I.; Alcott, Barry; Allen, Steven L.; Burrous, Jim; Cantrall, T.; Clarke, De; Coil, Alison L.; Cowley, David J.; Davis, Marc; Deich, William T. S.; Dietsch, Ken; Gilmore, David K.; Harper, Carol A.; Hilyard, David F.; Lewis, Jeffrey P.; McVeigh, Molly; Newman, Jeffrey; Osborne, Jack; Schiavon, Ricardo; Stover, Richard J.; Tucker, Dean; Wallace, Vernon; Wei, Mingzhi; Wirth, Gregory; Wright, Christopher A., "The DEIMOS spectrograph for the Keck II Telescope: integration and testing," Proc. SPIE, 4841, 1657 (2003)

[2] Oke, J. B., Cohen, J. G., Carr, M., Cromer, J., Dingizian, A. & Harris, F. H.,"The Keck Low-Resolution Imaging Spectrometer," PASP, 107, 375O (1995)

[3] McCarthy, J. K., Cohen, J. G., Butcher, B., Cromer, J., Croner, E., Douglas, W. R., Goeden, R. M., Grewal, T., Lu, B., Petrie, H. L., Weng, T., Weber, B., Koch, D. G., & Rodgers, J. M., "Blue Channel of the Keck low-resolution imaging spectrometer," Proc. SPIE, 3355, 81 (1998)

[4] Steidel, C. C., Shapley, A. E., Pettini, M., Adelberger, K. L., Erb, D. K., Reddy, N. A., Matthew, P., "A Survey of Star-forming Galaxies in the 1.4<~Z<~ 2.5 Redshift Desert: Overview," ApJ, 604, 534 (1994)

[5] Rockosi, C., Stover, R., Kibrick, R., Lockwood, C., Peck, M., Cowley, D., Bolte, M., Adkins, S., Alcott, B., Allen, S. L., Brown, B., Cabak, G., Hilyard, D., Kassis, M., Lanclos, K., Lewis, J., Pfister, T., Phillips, A., Robinson, L., Saylor, M., Thompson, M., Ward, J., Wei, M., Wright, C., "The low-resolution imaging spectrograph red channel CCD upgrade: fully depleted, high-resistivity CCDs for Keck," Proc. SPIE, 7735, 26 (2010)

[6] McLean, I. S., Steidel, C. C., Epps, H. W., Matthews, K. Y., Adkins, S. M., "MOSFIRE: the multi-object spectrometer for infrared exploration at the Keck Observatory," Proc SPIE (submitted) (2012)

[7] Adkins, S. M., Cohen, J. G., Aycock, J., Bell, J., Cohen, R., Cooper, A., Goodrich, R., Johnson, J., Kwok, S. H., Lyke, J., McCann, K., Neyman, C., Nordin, T., Panteleev, S., Tolleth, G. & Tsubota, M., "MAGIQ at the W. M. Keck Observatory: Initial deployment of a new acquisition, guiding, and image quality monitoring system," Proc. SPIE ,7014, 62 (2008)

[8] Kwok, S. H., Johnson, J., Adkins, S. M., McCann, K., "The Software for MAGIQ: a new acquisition, guiding and image quality monitoring system at the W. M. Keck Observatory," Proc. SPIE 7019, 9 (2008)

The W. M. Keck Observatory is operated as a scientific partnership among the California Institute of Technology, the University of California, and the National Aeronautics and Space Administration. The Observatory was made possible by the generous financial support of the W. M. Keck Foundation. The Observatory would like to acknowledge Gordon and Betty Moore for their generous support of the MOSFIRE instrument. This material is based in part upon work supported by AURA through the National Science Foundation under AURA Cooperative Agreement AST-0132798, as amended. The authors wish to recognize and acknowledge the very significant cultural role and reverence that the summit of Mauna Kea has always had within the indigenous Hawaiian community. We are most fortunate to have the opportunity to conduct observations from this mountain.