

The Elixir System: Data Characterization and Calibration at the Canada-France-Hawaii Telescope

E. A. MAGNIER¹ AND J.-C. CUILLANDRE²

Canada-France-Hawaii Telescope Corporation, 65-1238 Mamalahoa Highway, Kamuela, HI 96743

Received 2002 December 11; accepted 2004 February 26; published 2004 April 20

ABSTRACT. The Elixir System at the Canada-France-Hawaii Telescope performs data characterization and calibration for all data from the wide-field mosaic imagers CFH12K and MegaPrime. The project has several related goals, including monitoring data quality, providing high-quality master detrend images, determining the photometric and astrometric calibrations, and automatic preprocessing of images for queued service observing (QSO). The Elixir system has been used for all data obtained with CFH12K since the QSO project began in 2001 January. In addition, it has been used to process archival data from the CFH12K and all MegaPrime observations beginning in 2002 December. The Elixir system has been extremely successful in providing well-characterized data to the end observers, who may otherwise be overwhelmed by data-processing concerns.

1. INTRODUCTION

With recent advances in the large mosaic cameras, data-handling tasks that were once trivial have become onerous for most end users. Simple steps such as creating quality flat-field images have become a difficult job for users without significant investments in computer hardware. Even the bookkeeping required to keep track of the large number of objects that can be detected on a single set of images has become a significant database problem. The resulting intimidation many users feel inhibits some from requesting time with the large cameras, or acts as a barrier to the data-reduction task if they actually acquire the data. The groups that have been the most successful in producing results from the large mosaic cameras at the Canada-France-Hawaii Telescope (CFHT) and elsewhere have usually spent large amounts of time writing dedicated software for the job, and may even bring their own hardware to the telescope to analyze the data as it comes off the telescope to avoid the large overhead of writing to and reading from tapes. If these large mosaic cameras are to be used by the majority of observers, support must be provided to help them overcome these hurdles.

At the same time, with the start of queued service observing (QSO) at CFHT and many observatories, it is necessary for the observatories to take responsibility for data-reduction tasks that were once left to the observers. With data obtained over many different nights in queue mode, it is impractical for each observer to generate and interpret the large amount of calibra-

tion data that might be needed for even a few images. In addition, there is an advantage that can be achieved by having the observatories in charge of data calibration: they are best suited to monitor changes to the instruments and to determine improved calibrations based on very large samples of data. Finally, with the increasingly important role of archiving in astronomy, it is vital to provide a trustworthy calibration with the archived images, and even to apply the calibrations, if possible, for the archive users.

At CFHT, the Elixir project has the goal of enabling the optimal use of data by our observers, in particular from the large mosaic cameras. We call the project “Elixir,” after the fabled goal of the ancient alchemists: the Elixir that could restore youth or turn lead into gold. The Elixir project seeks to convert the lead of raw data into shining gold of well-calibrated images, restore youth to archived data, and turn the weighty mass of unmined data collected at CFHT into nuggets of gold.

There are several ways in which a dedicated project, having access to all acquired data, can smooth the data-reduction process:

1. All available detrend³ data can be used to produce optimal “master” detrend frames for a given set of science data. Having access to all detrend images during a particular camera “run” (the period during which the camera is mounted on the tele-

¹ University of Hawaii, Institute for Astronomy, 2680 Woodlawn Drive, Honolulu, HI 96822; eugene@cfht.hawaii.edu.

² Observatoire de Paris, 61 Avenue de l’Observatoire, 75014 Paris, France; jcc@cfht.hawaii.edu.

³ We are using the term “detrend” data to refer to the images needed to remove the purely instrumental signature—flat field, bias, etc.—as opposed to “calibration” data, which is used to measure externally relevant quantities—photometric zero points and astrometric terms. We distinguish between these two types of data because they are treated completely differently in all analyses. Analogously, “detrending” is the process of applying the detrend data to raw science images, resulting in “detrended” data.

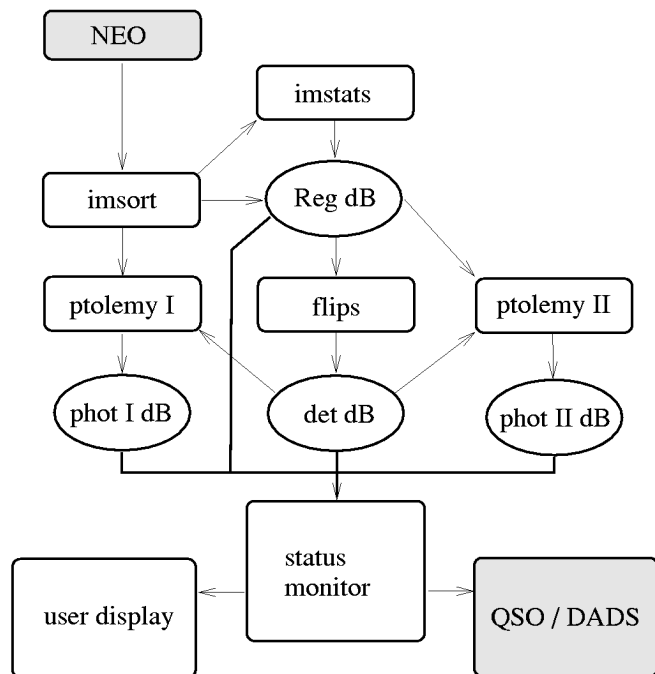


FIG. 1.—Conceptual diagram of data flow in Elixir. The arrows show the path that data about an image follow through the components of Elixir. Rounded rectangles represent programs, ellipses represent data products.

scope) avoids the limitation of using only the few images one can obtain in a single night or two at the telescope.

2. Constant monitoring of the photometric, astrometric, and other system-wide characteristics of the combined camera-filter-telescope system over all relevant timescales can aid in the identification of systemic problems and can allow for improved determination of calibration terms such as photometric zero points. We can determine on-the-fly whether a particular set of standard star images are taken in photometric conditions, and we can watch for slow trends if, for example, components acquire deposits or the detector quality fades as a result of contamination.

3. All data can be passed through a standard reduction system, providing observers a “quick look” mechanism in real-time and sanity checks when reducing the data. Dedicating computers and software to this task allows for optimization of the analysis routines to provide high-quality reductions on short timescales. This allows observers to decide in real time if their images are deep enough, if they cover the correct part of the sky, or if they have contamination from, e.g., bright stars. For some projects, the standard reduction results produced by Elixir may be sufficient for the intended science.

The Elixir project at CFHT is part of the encompassing New Observing Plan (NOP), which aims to increase the efficiency of observing at a variety of levels (Martin et al. 2000; Fahlman & Grundseth 2001). In addition to Elixir, the NOP consists of the QSO system, which has the goal of optimizing the obser-

vations performed at the telescope based on the actual weather conditions (Martin 2001); the New Environment for Observations (NEO), which serves to improve the efficiency of the observing process; and the Data Archive and Distribution System (DADS), which is responsible for storing the raw data and distributing raw and processed data to observers (Withington & Grundseth 2001). The role of Elixir in the NOP is similar to the role of the *Hubble Space Telescope* on-the-fly reprocessing (OTFR) system (Swade et al. 2001), since it provides observers with processed and calibrated images.

Since the NOP is a project still undergoing development, the various components of the NOP are currently applied only to specific instruments. In the case of Elixir, the CFH12K and MegaCam wide-field imaging cameras are the primary targets. The CFH12K camera consists of 12 $2K \times 4K$ MIT/Lincoln Labs CCDs arranged in a 2×6 pattern covering a total area of 0.3 deg^2 (Cuillandre et al. 2000; Starr et al. 2000). This camera was operational from 1999 February until 2003 January and was operated largely in QSO mode after 2001 January. MegaCam is the next-generation wide-field imager at CFHT and is part of a completely new prime-focus upper end called MegaPrime (Veillet 2000). MegaCam consists of 40 $2K \times 4.5K$ Marconi/EEV CCDs covering a roughly square region of approximately 1 deg^2 . MegaPrime achieved first light in 2002 December and has been used for regular science operations since the beginning of 2003. In addition to these two wide-field optical imagers, the Elixir system is used in a reduced capacity for the infrared imager CFHT-IR and will be extended to complete functionality for the future wide-field infrared imager WIRCam. CFHT-IR is a $1K \times 1K$ HAWAII array covering an area of roughly 13 arcmin^2 (Forveille 2001a). WIRCam will employ four $2K \times 2K$ HAWAII-2 arrays to extend the available imaging area to 400 arcmin^2 (Forveille 2001b).

2. ELIXIR SYSTEM OVERVIEW

The Elixir system consists of many independent software components, reference data, and databases that are connected together to perform a variety of analyses. Figure 1 shows the connections between the major Elixir subsystems at the highest level. In this diagram, the arrows trace the motion of information about images, groups of images, or derived quantities as it moves through the Elixir system. The different timescales and the different computer hardware usage is not represented, only the conceptual interactions between subsystems.

In Figure 1, analysis subsystems are represented by rectangles, while database tables and collections of tables are represented by the ellipses. Systems that are external to Elixir are identified as shaded boxes. The Elixir analysis blocks represent fairly complex collections of programs, discussed in some detail below (§ 3). Here we provide a general overview of the analysis that is performed.

Data from the camera on the telescope (represented by NEO) go directly to a set of Elixir processes (*realtime*) that run on

a computer at the summit in the CFHT dome and that perform basic analyses to provide feedback to the observers. The data are also delivered from the camera to the CFHT DADS in Waimea, which maintains all of the raw data for all needs in Waimea. DADS in turn delivers references to the images to the Waimea Elixir subsystems by sending information to the block labeled *imsort*. The *imsort* subsystem inserts references to all raw images into a database table (*reg.db*) and also sends a reference to the images to two other subsystems. The first of these subsystems, *imstats*, provides a very minimal, fast set of analyses on all of these images, with the results updating the same table. With a database of all images obtained by the telescope, and along with their characteristics, the later processing stages can make intelligent choices of images to process. The other subsystem that receives data from *imsort* is the block labeled *ptolemy (RT)*, which provides reduced images and potentially extracted object lists and astrometric calibrations, to the real-time analysis systems that are installed in Waimea for the CFHT Legacy Survey (CFHT-LS). All of the subsystems discussed in this paragraph run on data in more-or-less real time during the camera run.

The next set of processing subsystems shown in Figure 1 represent the end-of-run calibration analysis. This starts with the detrend and fringe creation subsystems, *mkdetrend* and *mkfringe*, which are applied to the collection of detrend data taken during a CFH12K run. These are followed by *ptolemy (PR)*, which performs a detailed analysis of all science images, including flat-fielding, photometry, astrometry, and inclusion of the detected objects in a stellar photometry database system (*phot.db*).⁴ The analysis performed by *ptolemy (RT)* and *ptolemy (PR)* are essentially identical, with only minor differences. The *standards* subsystem provides an analysis of the photometric standards and adds the results to the photometric database. The result of these steps is a complete set of calibration data (detrend images, astrometric calibrations, and photometry zero points) for the observing run.

The remaining analysis subsystem listed in Figure 1 represents the processing performed on raw images when they are prepared for distribution to the end users. This subsystem is called by the DADS system and applies all of the calibration information obtained by the end-of-run subsystems. The distribution of data is performed on a third timescale. Data are distributed to the observers under several conditions. Normally, observers receive their data at the end of the semester, or when their observing program is completed, whichever comes first. Observers may request to have their data earlier; for example, at the end of each run in which the observations are made. Data for the CFHT-LS are distributed at the end of each run as well, assuming the end-of-run calibration stage is ready for that run. Before images are released for distribution, an Elixir process not visible in this diagram must validate that all of the

calibration data necessary for that image have been generated. Some observers who require very fast access to their data may request processed data even if the calibrations are not available. In this case, the images are processed in a “best effort” mode, in which case the most recent available detrend images (e.g., from the previous camera run) are applied. This is the level of processing applied in the real-time stage for the CFHT-LS survey analysis systems.

3. MAJOR ELIXIR COMPONENTS AND EXTERNAL PACKAGES

The Elixir system is designed to be flexible. The interfaces between components, such as the image detrending and the object detection stages, are clearly defined so that other programs can be substituted fairly easily. There are several steps that invoke programs developed outside of the Elixir system. In addition, there are several large programs that are part of the Elixir system but that deserve substantial discussion on their own. In this section we discuss the major software components of the Elixir system and the external software systems used by Elixir. For further detailed information on the workings of these packages, please see the included references.

3.1. Parallelization and Program Organization: *gcontrol*

Most data analysis performed by Elixir is well suited to a chunky parallelization. The unit of parallelization may either be a single CCD image, a single mosaic frame, or a group of related CCD images, such as all flat *I* data from chip 03. To facilitate this type of operation, we have used the locally developed program, *gcontrol*.⁵ This program allows one to organize and coordinate a complex set of operations on a number of images or image groups, using an arbitrary number of clustered computers to perform the operations in parallel. The *gcontrol* program provides a mechanism to pass information from one program to the next to manage the analysis process. Depending on the configuration, *gcontrol* can define a wide variety of analysis systems by grouping together different programs as needed. Each implementation is defined by a simple text-based configuration script. We refer to a particular instance of *gcontrol* by appending the configuration name. For example, the *gcontrol* implementation that performs the Elixir *imstats* subsystem analysis would be referred to as *gcontrol:imstats*.

The appropriate use of *gcontrol* is for situations in which there are a large number of identical input data items, each data item needs to have the same sequence of tasks applied, and there is a collection of computers on which to perform the tasks; for example, a collection of science images that need to be flattened, defringed, and have object detection performed. There are thus four major tasks for *gcontrol*: First, provide a mechanism to define the sequence of tasks for an arbitrary

⁴ The Elixir Photometry Database Web page is <http://www.cfht.hawaii.edu/Instruments/Elixir/photomDB.html>.

⁵ For information on *gcontrol*, see the CFHT Web page <http://www.cfht.hawaii.edu/Instruments/Elixir/Ohana/gcontrol.html>.

(abstract) input data item. Second, generate the specific commands (tasks) for each data item in a collection of input data items. Third, assign tasks to available computers. Fourth, monitor the tasks in progress and manage the staging of tasks as they finish. To provide specifics, we will use the example of the image analysis steps given above. The tasks are assumed to be commands given at the UNIX shell level, and connections between the remote computers are performed with either RSH or SSH.

Imagine we have a collection of N images, `image.000.fits` to `image.NNN.fits`. We want to perform the three analysis steps given above: flatten, defringe, object detection (note that these are simply example command names, not actual Elixir names). For each of these steps, we have a UNIX shell-level command for each of these steps:

```
flatten image.fit image.ftt
defringe image.ftt image.def
getobjects image.def image.lst,
```

and we have a collection of M computers on which to perform these analyses. For the sake of illustration, assume that the images are available on all machines with exactly the same name; perhaps the name includes an absolute path with implicit cross-mounting of the necessary disk resources.

The input data list to *gcontrol* consists of a sequence of lines, each with a fixed number of fields. The *gcontrol* script defines the commands with abstract names for the inputs, and *gcontrol* constructs the commands for each line in the input data list on the basis of these fields. In this example, the input data list may consist of the strings `image.000` through `image.NNN`, with the extension stripped off:

```
image.000
image.001
image.002
...
image.NNN.
```

The fields in each line of the input data list correspond to the variables `&0–&N` in the *gcontrol* script. The script would define the flatten command to be generated from the input data; e.g., `flatten %s.fit %.ftt`, where the values of `%s` are replaced with the value of `&0` for each line in the input data list.

Each command defined in the *gcontrol* script is associated with one input queue and two output queues: success and failure. The script defines how the output queue for each command is connected to the input queue for another command. There is a special set of input and output queues, called “global,” that define the initial starting point and the end disposition of items. For example, for most commands, the failure output queue, where data items are placed if the command failed, is linked to the global failure queue. The last command in the chain should have its success queue linked to the global success queue.

As it is run, *gcontrol* monitors the input and output queues

for each of the commands and the processing state of the machines currently executing commands. When a machine is found to be idle, *gcontrol* will attempt to find a pending input data item waiting on the input queue of a command, construct the specific command for that input data item and command, and send the command to the idle machine. When the machine is finished with the task, *gcontrol* moves the data item to the appropriate output queue and onto the next appropriate input queue. In this way, *gcontrol* uses the cluster of computers as it moves all of the input data items through each of the command steps and eventually to the global success or failure queues. In addition, *gcontrol* will attempt to detect whether a computer crashes or is halted, rerun commands as necessary, and attempt to reconnect to the machine if possible.

3.2. Image Reduction and Combination: FLIPS

The FLIPS⁶ collection of programs is designed to perform the basic image reduction operations related to the detrending process. In this collection there are tools, for example, to merge several input flats to create a master flat or a master bias frame. There are also tools to apply the resulting detrend images to science images. The FLIPS tools are designed to make these steps fast and efficient and to make the manipulation of CCD mosaics transparent. We use FLIPS programs to create the master bias, dark, flat, and fringe frames and to apply these frames to the science images. The FLIPS tools are distributed as part of the Elixir system.

The two principle FLIPS programs used in Elixir are *imred* and *imcombred*. The first of these performs the steps of bias subtraction, dark subtraction, flat-field correction, and image masking. The program performs all of these operations in a single pass and only operates on a portion of the image array at a time. The input to the program includes the choices of which steps to perform and what detrend images to use. The bias-correction options used by Elixir include both a bias image and an overscan correction. The overscan is fitted with a low-order polynomial. If a mask pixels is set, the image pixel is set to the value of 0. Flat-field images are divided, not multiplied, and the output image can be optionally rescaled on the basis of the input flat-field mode. In Elixir, we do not request rescaling, because all chips of the flat field are created with a common normalization that ensures the flattened image will have a single zero point for all chips. The output FITS images are written with 16 bit integers, with the BSCALE and BZERO values set to maintain the original dynamic range, and also set so that the numerical precision of the image 0.0 value is 0.001. This latter point involves making small adjustments to the value of BZERO to force the data value of 0 to have an integer byte value in the output data. In addition, a random value with a range of $\pm \frac{1}{2}$ of an output bit is added to the data values before

⁶ See the CFHT Web page <http://www.cfht.hawaii.edu/Instruments/Flips>.

integer truncation. This avoids the segmented structures seen in low-noise integer data.

The second FLIPS program used by Elixir, *imcombred*, is used to combine input detrend images into a single master and to generate and apply the master fringe frames. In generating a single master detrend frame from an input stack, *imcombred* allows for several options for the combination statistics. There is the choice of what method to use to reject the outliers in the input stack (e.g., stars in the input flat-field image). In Elixir, we choose the CCD-CLIP statistic, which uses the CCD noise characteristics to predict the expected pixel value standard deviation and reject pixels that lie more than a specified number of standard deviations (3 is used by Elixir) from the mode. The other choice is the statistic used for the pixel value. With Elixir, we use the median if there are more than six input images, and otherwise we use the mean. In the fringe master creation, the strength of the fringe signal on the input images is measured, and the input images are combined after a base sky level has been subtracted and the remaining fringe component is scaled by the measured fringe amplitude. The fringe and sky measurements are not performed by *imcombred*, but rather by the Elixir supporting components, and the results are included in the inputs to the *imcombred* task. In the fringe master application, the fringe amplitude is again measured, and the master is scaled to match before it is subtracted from the science image.

3.3. Object Detection and Classification: *SExtractor* and *GoPhot*

The *SExtractor* package (Bertin & Arnouts 1996) is an efficient and easy-to-use tool for performing stellar photometry on an image. The output is very flexible and can easily include as a wide variety of measurable quantities as necessary for each object detected. *SExtractor* can be run with very little initial information, which makes it particularly useful in an automatic-processing environment. The program can be obtained from the Web, although the version currently used by Elixir is included in the Elixir distribution. Within the Elixir system, we use *SExtractor* as our primary object detection and classification tool. Although we find that the object classification scheme is not as robust as some other packages, like *DoPhot* (see below), it is extremely fast. In most of the Elixir system, the speed requirements outweigh the requirement of good object classification. For the standard object detection used in the *ptolemy* subsystem and thereafter used by the standard star photometry system, we record several values for each object (using the names given by *SExtractor*): CLASS_STAR, X_IMAGE, Y_IMAGE, MAG_BEST, MAGERR_BEST, BACKGROUND, FWHM_IMAGE, A_IMAGE, THETA_IMAGE, MAG_ISO, MAG_APER, and FLAGS. For some of the other quick analyses, such as the focus analysis tool in the *realtime* subsystem or the seeing measurement, we only extract the positions, FWHM, magnitudes, and object flags.

An alternative object detection and classification tool avail-

able with Elixir is *GoPhot*. *GoPhot* is our adaptation of the program *DoPhot* (Mateo & Schechter 1989). The *DoPhot* algorithm measures stellar photometry of objects in an image by performing analytical fits to the object profiles. Unlike *SExtractor*, *DoPhot* uses a simple two-dimensional Gaussian profile to measure the photometry. *DoPhot* has a somewhat more physical classification scheme for the detected objects than *SExtractor*, but it is somewhat slower. *DoPhot* performs object detection in a series of stages, fitting and subtracting the brightest objects in one stage, then revisiting the entire image again at a lower brightness threshold, stopping when a predefined threshold is reached. *GoPhot* is our conversion of the *DoPhot* code to C, along with minor improvements, including better handling of saturated stars and large diffuse objects. *GoPhot* (and *DoPhot*) can be significantly slower than *SExtractor*, because the routine requires fitting a Gaussian profile to every object several times as the reference stellar profile is improved.

3.4. Other Elixir Components

MANA⁷ is a command-line-driven image analysis package that includes an extensive interpretive programming language. This program includes the tools needed to manipulate and display one-dimensional (vector) and two-dimensional (image) data, as well as extensive arithmetic operations. This tool is used by the Elixir system for vector math in several subsystems and for creating useful displays of images and plots of data, and is distributed with the Elixir system.

The Apache⁸ HTTP server is an open-source Web server, and the most popular Web server at present. It is robust, secure, and is maintained by an active community that keeps on top of changes in the security requirements. The Apache HTTP server is used by the Elixir system to serve the Web pages used to evaluate the quality of the detrend images. Since it is used purely for internal purposes, it does not need to accept outside connections, making the server even more secure. Although the Elixir system uses Apache, any Web server would function appropriately.

Most programs in the Elixir system do not require human interaction and therefore run “in the background” without a terminal or windowing system necessary. However, certain programs such as MANA generate graphical output using X Window System tools. It is convenient to have a guaranteed X server available to these programs, without the uncertainty of access and availability of the console of a given computer. The Elixir system uses the program *Xvnc*⁹ to provide a virtual X server for which the availability is more easily controlled. In the arrangement at CFHT, we use two separate *Xvnc* servers—one in the summit network and one in the network at the headquarters in Waimea.

⁷ A MANA user’s guide, <http://www.cfht.hawaii.edu/Instruments/Elixir/Ohana/mana.html>.

⁸ The Apache server can be found at <http://www.apache.org>.

⁹ Available at <http://www.uk.research.att.com/vnc/xvnc.html>.

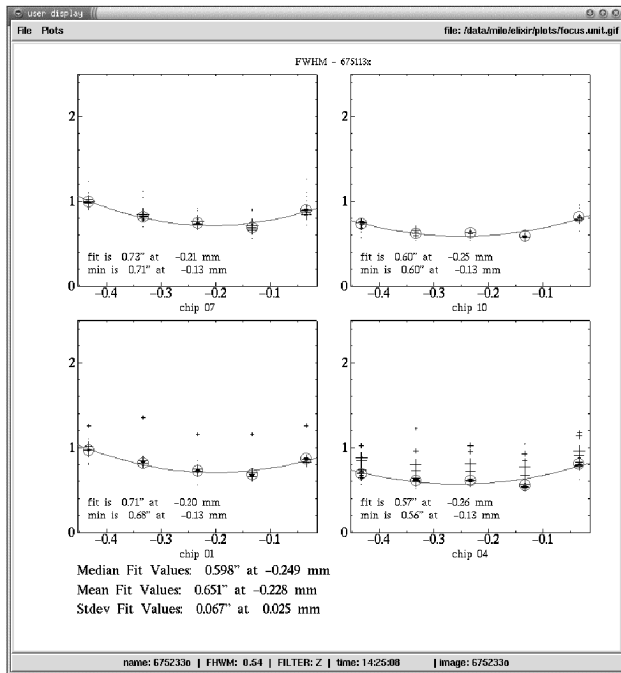


FIG. 2.—Sample focus analysis plot. The four panels represent measurements from four CCDs symmetrically spaced about the optical axis. Each group of crosses represents the FWHM of the stellar images from the given focus value. The circles represent the median value of these groups, and the parabola is a fit to the circles.

4. ELIXIR SUBSYSTEMS

In this section, we provide additional information on each of the major subsystems discussed in § 2. The Elixir subsystems make use of the various components discussed in § 3.

4.1. *Realtime*

The *realtime* subsystem consists of several independent processes that provide immediate feedback to the observers. While the rest of the analysis systems run on machines in Waimea and receive their data from DADS, the *realtime* processes are triggered directly by the data acquisition system. In addition, the analysis of each image is synchronous in the sense that the processes are immediately launched for each image, and therefore are completed in a generally predictable time from the acquisition of an image. This also implies that the images are not buffered for this analysis; if the analysis of an image takes too long before the next image arrives, the processing is aborted so that the most recent image is analyzed.

The analyses performed include: (1) seeing measurement, (2) creation of a binned, gray-scale jpeg image, and (3) analysis of a focus frame. The results from these *realtime* processes, along with other asynchronous results discussed below, can be viewed by the observer within a single display tool. The seeing measurement is basically identical to that performed by *imstats*,

discussed below, but only a single chip near the field center is analyzed.

For the focus analysis, the focus images are obtained with exposures at multiple focus positions integrated on a single frame, with the telescope (or detector charge) offset by a fixed amount between each exposure ($2 \times$ on the last). The focus analysis uses SExtractor to measure the FWHM of all objects in the image and then identifies the object groups and determines the sequence by keying on the double-spaced pair. As a result, the FWHM values for all stellar images obtained at each of the focus positions can be accumulated, and a statistic (i.e., the median) determined, and the focus curve fitted. This process is performed on four of the detectors from each focus frame so that the best focus is chosen for an annulus 50% of the mosaic radius. An example of the focus analysis plot from CFH12K is shown in Figure 2.

As of fall 2002, the *realtime* analyses are performed on a dual-CPU 1.2 GHz Intel Pentium computer running an in-house distribution of Linux. On this computer, the focus analysis requires roughly 8 s, which means the plot is available to the observers nearly as quickly as they can view the focus image, much less analyze the image shapes. In the case of the seeing measurements, the analysis is performed in under 2 s, including the time for the graphical display to update to the most recent value. Speed is particularly crucial for the QSO system to minimize overhead in the decision-making process.

4.2. *Imstats*

The quick-statistics subsystem *imstats* performs a few basic measurements on each CCD image and places the results in a database of registered images. For the optical wide-field imagers, a first component measures the bias level and sky brightness, while a second element uses SExtractor to measure the FWHM of the brightest stars, down to 7σ above the background. To speed up this analysis, a small segment of each chip, limited to 1600×1600 pixels, is used. This size was chosen to balance the need for speed with the need for a robust measurement based on a sufficient number of stars. This region corresponds to roughly $5'.3$, and generally contains a sufficient number of stars to provide a reliable seeing measurement. The extracted collection of stellar measurements is somewhat filtered to provide the single seeing value. First, stars that are likely to be saturated or that are otherwise flagged with an error flag by SExtractor are excluded. Next, the peak of the FWHM distribution is found. Only measurements within 0.2 pixels of the peak are kept, and the mean of their FWHM values is taken as the image FWHM. We find that these measurements of the image quality are generally consistent with more detailed measurements by hand at a level of $0''.05$. For the infrared imager CFHT-IR, the entire field is used, since the field is much smaller and exposures are generally much more shallow. In addition, the depth is increased to 3σ in this case. Note that these seeing measurements are performed on images without flattening, al-

though in the case of CFHT-IR, a low-order polynomial fit to the sky background is subtracted to enhance the detection of stars. Note that there is no system feedback to check the quality of the seeing measurement; if there are too few objects, or if a chip is located on a galaxy cluster with few unresolved objects, the result may be somewhat biased. For the case of the full mosaic, only a fraction of the images will suffer this type of failure, so more reliable information can be obtained by examining several chips. Our experience is that only rarely do images result in seeing measurements that differ substantially from a more detailed examination by an observer.

The entire *imstats* operation is meant to happen reasonably quickly after the image has been taken so that the observers can have near-real-time feedback via the user display tool, also used to display the output from the *realtime* components. The processing takes place on machines in the Waimea Linux cluster, where a range of computing resources are available. We find that the complete *imstats* process for the optical cameras requires roughly 500 clock cycles per pixel, or roughly 4.5 s for a single CFH12K CCD on a 1 GHz computer. Since the *imstats* processing is typically distributed on six computers, the system is able to keep up with most typical data rates. The user display tool includes plots of the FWHM and sky brightness as a function of time for a recent time period. In addition, the seeing measurements for the full night, and the most recent 3 hr, are displayed on the Web for reference by other Mauna Kea observatories. An Elixir system that runs in the background updates these plots as needed by extracting the relevant data from the Image Registration Database.

4.3. *Mkdetrend* and *mkfringe*

The detrend creation portion of the Elixir system is divided into two stages: *mkdetrend* and *mkfringe*. Currently, these tasks are only performed for the optical wide-field imagers, and not CFHT-IR. The first of these processes, *mkdetrend*, is responsible for generating the master bias, dark, and flat-field frames from the raw images. Once these first-level detrend frames have been created, it is then possible to generate the additive correction frames, including both fringe frames and frames to correct the large-scale additive structures. This latter task falls to the system *mkfringe*.

Example processing of the raw detrend frames is automatically performed during a camera run, the period in which the imager is mounted on the telescope and data are being collected. This is used by observers to decide whether good flat-field images can be constructed or if better input images are needed. However, in general, the Elixir system defers the creation of the final master detrend images until the camera run has completed. The camera run defines a timescale over which the detrend data are likely to be stable. We therefore use this timescale as a starting point, and in the process of master detrend creation we test the consistency of the detrend images for the camera run. We have found that in general a single set of bias,

dark, flat, and fringe frames can be applied to an entire camera run. To date, there have been only two occasions when we have found it necessary to divide the camera run into different periods because the flat-field images changed significantly. In these instances, the removal of the CFH12K shutter allowed dust particles to fall on the exposed filter below.

We have made some useful advances in the handling of the flat fields and the correction for additive structures in the images. Regarding the flat-field effects, we have found that the flat-field images for the CFH12K camera require correction in order to be photometrically flat. This is for two likely reasons: first, the geometric correction introduced by the optical distortion in the camera, which is well-known and can in principal be corrected analytically; and second, the effect of scattered light, which contributes extra light across the focal plane, but can be modified approximately by a vignetting pattern. The effect of both contributions is to elevate the flat in the middle of the detector and to depress the flat near the corners. We have found that the simplest and most direct way to correct these effects is to use a grid of offset images taken in photometric conditions to measure the introduced error. This error can then be converted to an image that can be applied to the original twilight-flat images. We have found that this photometric correction is quite stable over long periods of time, and a single correction has been applied to all CFH12K data obtained to date. This flat-field process results in relative photometry across the mosaic that is consistent to 0.7%–1.0%. The issues involved are discussed in further detail in § 5.

In the realm of fringe correction, we have developed a method of correcting both the fringe pattern, which varies on high spatial frequencies but relatively low temporal frequencies, as well as a variety of other additive components with very large spatial scales, but which may change significantly from image to image. We have found that by independently treating the high spatial frequency component of the fringe pattern and the low spatial frequency components, a single fringe master can be applied successfully to all images from a camera run period. The residual in the fringe frames that we achieve is typically in the range of 5–10 counts peak-to-peak on a background of 3000–5000 in the *I* band. The low-frequency structures result from several sources, including scattered moonlight, varying filter response as a function of incident angle, differences between the spectral energy emission of the nighttime and twilight sky, etc. Since these terms can vary significantly and independently, we have found that it is necessary to decompose the background of a given image into principal components to adequately correct these effects with finite computing resources. We have used single-value decomposition to construct an appropriate set of basis functions that describe these low-frequency structures. For CFH12K we used several hundred images in each of the relevant filters obtained over the course of 6 months to generate the basis function. Once these principal modes have been constructed, they can be applied to data spanning years of operation. For further

details on the additive components, see Magnier & Cuillandre (2004).

The *mkdetrend* and *mkfringe* systems provide the necessary organization to the process of master detrend frame creation. They use the image registration database to make the initial image selections, and then launch processing pipelines as needed to process and merge the input images. The operation of the *mkdetrend* and *mkfringe* systems currently requires some human intervention. When these systems generate their master detrend frames, they also produce residual images and statistics on the residuals of the input detrend images to aid in improving the selection of input images. A tool that uses Web forms for the interaction makes it easy to evaluate the selection of the images used to create the master detrend frames and alter the input as needed. Figure 3 illustrates this tool in action. Once the operator is satisfied by the resulting master detrend images, the *mkdetrend* system automatically registers them in a database of detrend images for use by other Elixir subsystems. These images are also automatically available on the CFHT Web site, for users in the outside world.¹⁰

4.4. Ptolemy

The detailed analysis system *ptolemy* performs a complete photometric and astrometric analysis of each CCD image: detrending, object detection, flux measurement, astrometric calibration, and incorporation into a photometry database. The *ptolemy* analysis provides the measurements needed to assess the standard star photometry as well as the astrometric information for images to be distributed. All science images obtained during a camera run are passed through *ptolemy* once the master detrend images have been generated. The photometric detection is performed to a depth of 5σ above the background.

A variant of the *ptolemy* system is also run in real time on all images as they are obtained, using the best available detrend images for the task, which are likely to be generated from the previous camera run. To increase the speed, the photometric analysis is performed only to a moderate depth of 15σ , without pushing for the detection of the faintest stars in the image. This component is used to provide detrended images and data products to the CFHT Legacy Survey real-time analysis systems installed at CFHT by external scientific collaborations. These real-time systems can subscribe to any of the *ptolemy* data products, including flattened images, defringed images, output SExtractor object lists, and astrometric solution files. The requested data products are pushed to the data volumes registered with Elixir by the real-time systems.

Each of these steps is performed by a separate program, thus having a modular nature that allows for substitution of different components as needed. For example, the system currently can use either the program DoPhot or SExtractor to perform the

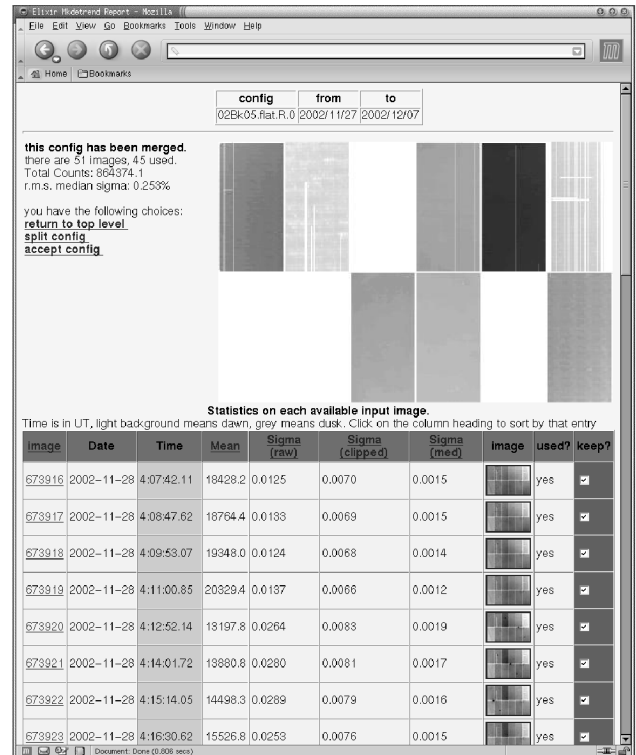


FIG. 3.—Example of the *mkdetrend* user interface tool. The tool allows the Elixir team to refine the selection of input images used to create a master detrend frame. The large gray scale shows the current master frame, while the information below provide statistics and thumbnail images of the residuals of the input images. The user selects or excludes images by checking the buttons, and resubmits the detrend data from processing after new selections have been made.

object detection/flux measurement step. Other programs may be easily substituted if necessary. We find that the complete *ptolemy* process typically requires ~ 4800 clock cycles per pixel, or about 40 s per chip on a single 1 GHz computer. By the end of the CFH12K period, the total processing resources were roughly 13 GHz, so a typical camera run can be processed through the *ptolemy* system in 5–10 hr, or somewhat longer if the network bandwidth is being consumed by other tasks as well. Six months after MegaPrime was introduced, the total processing power available for this process had increased to about 35 GHz, but the total data volume for each run had increased by more than a factor of 4. As a result, the typical *ptolemy* run in late 2003 for MegaPrime took between 20 and 30 hr.

4.5. Other Subsystems

The photometry database is used by the *standards* component to determine the photometric calibration parameters. The photometry database includes high-quality photometric standards from Landolt (1992) and others as needed. Queries to

¹⁰ The images can be found at <http://www.cfht.hawaii.edu/Instruments/Elixir/detrend>.

TABLE 1
ELIXIR & DADS COMPUTER SUMMARY

System	Quantity
Elixir data host computers	5 × 2 CPU (500, 750 MHz)
DADS (archive) data host computers	4 × 2 CPU (500, 750 MHz)
Elixir data space	2.1 TB
DADS data space	5 TB
Common processing computers (MHz)	6 × 1300

this database by the *standards* system are used to extract the CFH12K observations of these standards for the different filters, and to determine the zero point for each image. These measured zero points are included in a database table that gives the zero-point history. This table is used to generate statistics for each night, including the measured average zero point for a given filter for a given camera run, the comparison with the long-term average under photometric conditions, and the scatter for the camera run, after nonphotometric images have been rejected. These statistics indicate the reliability of the photometric solution and are included in the headers of images processed by the distribution system. We find that the zero point is typically consistent to a level of better than 1% over the course of a camera run if nonphotometric images are excluded (Magnier 2004).

The elements on the bottom of Figure 1 indicate the data visualization tools and the interaction with the rest of the telescope observing environment. Several tools exist for querying the Elixir databases. There are tools that can be used to select subsets of the data in the detrend and image registration databases, or to explore the photometry database. Such queries can be used to generate summary plots for the observers or for inclusion in the distribution package.

Data obtained by the CFHT QSO team are distributed to observers by DADS, which makes use of an Elixir component to perform the image processing. All images are detrended, and both the improved Elixir astrometric solutions and the photometric calibrations determined by Elixir from the standard star data for the run are added to the images headers. In addition, Elixir generates gray-scale thumbnail jpeg images, which are used by DADS to create a very useful data manifest in the form of a CD ROM that can be viewed with a Web browser.

An important Elixir subsystem is the SkyProbe atmosphere transparency measurement system. This system, described in detail in Cuillandre et al. (2002), consists of a 768×512 pixel CCD and 50 mm camera lens mounted on the telescope, with optical axes roughly co-aligned. The camera observes a $\approx 5^\circ \times 7^\circ$ region every 60 s. These images, which are sensitive to about 11th magnitude, are analyzed by an Elixir system that is very similar to the *ptolemy* analysis system described above. The resulting stellar photometric measurements are compared with the Tycho catalog of bright stars (Høg et al. 2000), and an image zero point is calculated. The difference between the

observed zero point and the nominal zero point gives the atmospheric transparency and is plotted for the observer.

While the bulk of the development effort has gone into the Elixir software, a vital element in the Elixir system is the computer hardware that is necessary to run the system. The Elixir computing infrastructure has been growing over the past 2 yr at CFHT, partly to improve the speed and organization of the system, but also to make the system ready for the deluge of data expected when MegaCam begins full operation.

The Elixir system can run on any standard UNIX or UNIX-like system. At CFHT, we are using a cluster of Linux computers, mostly Pentium III and IV systems from Dell. The Elixir and DADS projects at CFHT have some overlapping resources. The current system consists of several machines for both processing and Elixir data storage: several machines used primarily for the DADS data storage, and a group of machines dedicated to processing. These machines are on a network that is separate from the rest of the CFHT Waimea machines, and are connected with a 100 MB/1 GB switch. The Elixir parallel processing system is quite flexible about how many and which machines it uses for a given task. The particular allocation of machines varies, depending on the current demands and conflicting needs of the Elixir and DADS systems. The Elixir data machines store the Elixir reference data (i.e., the USNO catalog, configuration information, etc.), the master detrend data, photometry and image databases, and processing results from both the *ptolemy* and *mkdetrend* analysis systems. The DADS data computers are responsible for storing all raw images. Table 1 lists the computer hardware used by Elixir and DADS as of mid-2002. With the arrival of MegaCam, the hardware resources will be expanded to cope with the significantly higher volume of data.

5. FLAT-FIELD DETAILS

The construction of an appropriate flat field for camera systems such as CFH12K and MegaPrime is made more complex by the wide field of view. In the Elixir system, our strategy is to generate a flat-field image for the complete mosaic that brings all areas of the mosaic (all portions of all chips) to a common photometric system. This means that the flat field should result in individual CCDs that all have the same zero point. In this section, we discuss some of the difficulties involved in flat-field construction, and the choices we have made at CFHT. We start by justifying our preference for twilight flats within the Elixir system, and for NOP in general. We then discuss systematic errors observed in flat-field images: our explanation and justification for the source of the systematic error, and our strategies for correcting the error.

5.1. Use of Twilight Flats

The starting point for all flat-field images is some uniform illumination source. The traditional choices for the source include the twilight sky (“twilight flat”), a region in the interior

of the dome that is uniformly illuminated (“dome flat”), and the nighttime sky, combining many images to remove the contamination from astronomical sources (often called a “super-flat”). Within the Elixir system, we have elected to use the twilight sky as our illumination source. Each of the three types of illumination sources has advantages and disadvantages. We briefly present our rationale for choosing twilight flats.

The three main concerns of the illumination source are: (1) the illumination source must have sufficient spatial uniformity, (2) the spectral energy distribution of the illumination source must be sufficiently similar to that of the objects of interest, and (3) observations of the illumination source must be reliable enough that sufficient signal can be obtained.

We have avoided the use of nighttime superflat images primarily because of the problem of fringes. The night-sky spectral energy distribution is strongly dominated by line emission, mostly molecular oxygen and water vapor lines. The line emission is especially a concern in the long-wavelength filters, because thin-film interference within the detector causes fringe patterns to appear in the night-sky image. Since the fringe pattern is extremely different (washed out and reduced in amplitude) under continuum emission, the pattern is inappropriate for photometry observations of most astronomical sources, which are dominated by continuum emission. In addition, it is difficult, especially under the observing conditions of QSO, to obtain sufficient observations of the nighttime sky to generate superflats with significant signal-to-noise ratio.

We have also avoided dome flats within Elixir mostly out of concern about the uniform illumination pattern. Because of the difficulties involved at CFHT of setting up the illumination of the dome-flat screen, we have not explored the issue in as much detail as the night-sky superflat problem. Our original concern was that the flat-field screen illumination was not sufficiently uniform to produce acceptable flat-field images.

Twilight flats have the twin advantages of having a continuum spectral energy distribution, dominated by Rayleigh-scattering of sunlight, and an extremely uniform illumination pattern if the sky is photometric. Under photometric conditions, the twilight sky is extremely uniform, but small amounts of cirrus can introduce significant spatial variations. The possible large-scale gradients in the illumination are not a cause for concern, because of the corrections we discuss below. The main difficulty in using twilight flats is in obtaining the observations given (1) the short period over which the sky is usefully bright, and (2) the frequency of cirrus in the twilight sky. At CFHT, we find that observations performed within the QSO system overcome these obstacles by (1) providing the observers with sufficient tools and experience to catch the flat-field period, and (2) carefully monitoring the sky conditions to avoid periods of significant cirrus clouds. At Mauna Kea, there seem to be a sufficient number of photometric nights that we have been able to obtain the twilight flats needed for each run.

5.2. Flat-Field Systematic Errors: Causes

We find that regardless of the source of the flat-field images, there are systematic errors in the flat-field structure that we conclude are caused by scattered light contaminating the focal plane. Such an effect has been described for the ESO Wide-Field Imager (Manfroid et al. 2001). An example of the systematic error for CFH12K can be seen in the top panel of Figure 4, which shows the residuals of *R*-band standard-star photometry from the first QSO observing run with CFH12K as a function of mosaic *X* coordinate. This run provides an excellent example of the systematic error, because the photometric conditions of the sky were exceptional during the entire run; and because it was the first QSO run, the entire NOP team was paying extra attention to all of the factors that could have affected data quality. The amplitude of the systematic errors observed are in the range of 5% peak-to-peak.

We note that because of optical distortion in the camera, a flat-field image created on the basis of an illumination source with uniform surface brightness will introduce a similar type of error when used for stellar photometry. This effect has been discussed extensively (see, e.g., Manfroid et al. 2001), and we summarize the concept. As a result of distortion, the subtended surface area of a pixel at large field angle is smaller (in CFH12K) than that near the center of the mosaic. Since the flat-field source (i.e., the twilight sky) has a constant surface brightness, these pixels receive a smaller total flux than those at the center of the field. In a science image that is corrected with such a flat, the night-sky background, which is generally relatively uniform, will be corrected in exactly the same way and will appear flat. However, stellar photometry depends on the total flux, not the surface brightness. Therefore, in such a science image, the stellar photometry will be enhanced at the corners relative to the center of the mosaic.

However, the distortion error is a small contribution to the systematic error shown in Figure 4. The amplitude of the distortion error is small ($\sim 2\%$) compared with the observed systematic trend: the solid line in the top plot of Figure 4 shows the amplitude of the deviation caused by the varying effective pixel area in CFH12K; the amplitude is much smaller than the amplitude of the observed systematic error. In addition, as we show below, the observed systematic error as a function of position in the mosaic differs significantly from filter to filter, while the distortion should be largely achromatic.

We have explored possible causes for the observed effect. We find that the systematic error does not depend on the source of the flat-field image; twilight, dome, and night-sky flats all contribute similar, although not necessarily identical, errors. Different photometry analysis programs (e.g., SExtractor or DoPhot) result in the same systematic errors. The choice of the standard-star field does not affect the result.

Scattered light is an obvious culprit for this effect. If light is reaching the detector from sources other than via reflection from the primary mirror, it is likely that the resulting image




FIG. 4.—Standard star residuals resulting from three flat-field iterations. *Left*: Residuals as a function of X coordinate on the full mosaic field (*crosses*) for the simple twilight flat (*top*), the first attempt at a correction (photflat A, *middle*), and the final correction based on dithered images (photflat B, *bottom*). The top figure also shows the amplitude of the geometric distortion effect on the residuals. *Right*: Histograms of the same three residual sets. The bottom of these three histograms includes a Gaussian with $\sigma = 0.02$ mag that well represents the wings of the distribution but is too wide for the core.

will not adequately correct the detector response. If the effect is caused by scattered light, it is interesting to note then that the contamination seems to be very consistent for a wide range of twilight sky brightness values. This is seen in the fact that flat-field images obtained during twilight naturally span a large range of sky brightness values, necessitating exposures ranging from near the short limit (1 s) to nearly 100 s. Despite the large dynamic range in the sky brightness, the flat-field images are extremely consistent, at the under 1% level. However, the night-sky images are not so consistent; images with the moon above the horizon are significantly different from those taken without the moon. In addition, the closer the moon is to the optical axis, the more significant the deviation. Our conclusion from these clues is that the amplitude (and pattern) of the systematic error depends on the ratio of the light in the dome to the sky brightness; when the moon illuminates the inside of the dome during the nighttime, the observed sky image is substantially different from other periods.

We performed a test that is illustrative of this last point, that the pattern of light falling on the detector depends on the ratio of the dome light to the observed sky brightness. We obtained a series of twilight flats in photometric weather conditions with the dome slit severely constricted. To achieve this, we closed the shutter part way and raised the wind screen so that only a

small square region somewhat larger than the outer diameter of the upper ring remained open. We then pointed the telescope through this reduced aperture and obtained the twilight flats with this arrangement. In this layout, the primary mirror illuminates the detector as it normally would; the telescope beam is not vignetted by the dome slit. However, the interior of the dome is drastically darker than it normally would be for the same twilight sky brightness. The result was that the twilight flats obtained with this arrangement were substantially different from those obtained in the normal mode.

Under the assumption that the systematic errors are caused by scattered light, we attempted to identify possible sources of the scattered light. A detailed examination of possible reflecting light sources under realistic lighting conditions was performed by converting the CFH12K into a pinhole camera. We created a filter slide that could hold a thin sheet of metal in place of a filter, in which we placed 13 holes, 200 μm in diameter. Each hole acts as a pinhole camera, projecting on the detector an image of whatever is on the other side of the hole—in this case, the primary mirror and the support structures.

Figure 5 shows the pinhole camera images, including the full CFH12K mosaic field (*top*) and a zoomed version of the central primary mirror image (*bottom*). The 13 annuli scattered across the field are images of the primary mirror projected by




FIG. 5.—CFH12K pinhole camera images of the primary mirror and surrounding regions. *Top*: Positive image of the full field of view showing the 13 pinhole images. *Bottom*: Negative enlargement of the central image showing the mirror covers.




FIG. 6.—Scattered-light correction frames for CFH12K *R* flats. *Top*: Ad hoc correction (scatter A). *Bottom*: Correction measured on the basis of stellar photometry (scatter B).




FIG. 7.—Dither pattern used to measure the CFH12K flat-field correction (scatter B). The lines represent the outline of the individual CCDs in the mosaic, while the inset table gives the applied offsets in arcseconds for the R.A. and decl. sequences.

each of the 13 pinholes. In each image, the main circular structure is the primary mirror, with a dark shadow of the prime focus cage, as well as the spider legs of the support structures. Around the outside of the primary mirror, there are a series of trapezoidal shapes; these are the mirror covers, which are clearly the brightest sources other than the primary mirror.

The only obvious sources of light other than the primary mirror are the mirror cover petals. We arranged to directly measure the contribution of light reflected from the mirror cover petals. In 2001 April, we obtained a series of dome flats with the petals exposed, and again with large sheets of black cloth draped over the petals. We normalized the images and averaged them, then subtracted the “shroud on” from the “shroud off” images. The resulting image, Figure 6, consists of the excess illumination introduced by reflections off of the primary mirror cover petals. This figure shows the difference image for the *R* filter and clearly demonstrates the presence of excess light from the cover petals. The morphology of this image is roughly the appropriate shape needed to correct the photometry errors seen in Figure 4. However, the amplitude of the excess light term in these difference images is too small by a factor of roughly 10. This experiment implies that the mirror cover petals were not the principal source of the scattered light. We nonetheless removed the white Teflon pads that had the highest albedo, but there was no significant change in the flat-field pattern. As a result of this set of experiments, we conclude that the excess light reaching the detector comes from the general ambient light within the dome, scattered at a very low level off the many blackened surface visible to the detectors. The large am-

plitude of the scattered light contamination is due to the fact that the primary mirror subtends a small solid angle as seen from the focal plane, and the blackened surfaces that are contributing the scattered light subtend a very large angle. The low scattered light fraction is outweighed by the large ratio of surface areas.

5.3. Flat-Field Systematic Errors: Initial Ad Hoc Correction

Lacking any other correction options, we initially created an ad hoc correction for CFH12K based on the flat-field observations obtained with the mirror petals covered and exposed. Since the observed pattern was generally similar to the observed error, we used the pattern to correct the basic flat-field images. We used only the *R*-band contamination frame and applied it to each of the major broadband filter flat-field images (*BVRI*), adjusting the amplitude of the contamination frame to minimize the residuals of the standard star observations. The Elixir flat-field images are normalized so that a reference CCD has a median value of 1; in the case of CFH12K, the reference is CCD 04, while for MegaCam it is CCD 00. The application of the contamination frames described above therefore involved subtracting the contamination frames, multiplied by the determined scaling factor, from the flat-field images, and renormalizing the result so that CCD 04 retained a median of 1. In our online documentation, we call this correction frame “scatter-A.0,” and the corrected flat-field images receive the label “photflat-A.0.”

The reduction in the photometric residuals using scatter-A.0 was substantial. Figure 4 shows the residual plots for the standard star observations from the CFHT 2001A semester, with and without this correction, as well as the improved correction discussed below. In Figure 4, the left-hand plots show the standard star residual as a function of *X* coordinate in the mosaic, while the right-hand plots show the residual histograms for each of the three data sets. The middle pair shows the residuals when the scatter-A.0 correction is applied to the flat-field, while the bottom pair shows the improved correction discussed below. To determine these residuals, fixed linear air mass and color corrections were applied to the instrumental photometry, and a single zero-point offset was determined for each mosaic frame (not for each CCD independently).

5.4. Flat-Field Systematic Errors: Empirical Correction

Since it was clear we could not eliminate all sources of scattered light in the flat-field images, we decided to construct a correction frame by measuring the effect of the contamination on stellar photometry. Such a correction has the advantage of correcting the observed error of concern. To make such a correction, we obtained, in photometric weather, a number of images with large dithered offsets so that a given star would be observed at a wide variety of mosaic positions. The correction is generated by using these repeated observations of the same




FIG. 8.—Stellar residuals from the sequence of dithered images used to construct the CFH12K *R*-band flat-field correction (scatter B). The bottom left panel shows the stellar residuals without the correction, while the top left panel show the residuals after the correction is applied. The histograms at right show the residual distribution for all stars (*thin line*) and for stars with formal errors of 0.01 mag. The smooth curve overlapping this distribution is a Gaussian with $\sigma = 0.01$ mag (*heavy line*).

stars to generate a map of the photometric error as a function of position in the mosaic. We discuss the application of the technique to CFH12K data, but note that we have successfully performed the same operation for MegaPrime.

We obtained the necessary sets of dithered images for CFH12K in photometric weather for each of the main wideband filters (*BVR*) during the several QSO runs in late 2001 and early 2002. We later also obtained data for the z' -filter and specific other CFH12K filters. When MegaPrime became available, we also obtained the same type of observations. For each filter, a set of images are obtained at 12 pointings (13 for MegaPrime), with offsets ranging from 50 pixels to half of the mosaic size in each of the *X* and *Y* directions (see Fig. 7).

We flattened these images with the appropriate uncorrected twilight master flat-field images from the corresponding camera run. We only used data obtained in photometric conditions as demonstrated by SkyProbe (Cuillandre et al. 2002). We then performed SExtractor photometry on the images, performed astrometry, and included the measurements in the Elixir photometry databasing system.

We divided the entire mosaic area ($12,500 \times 8200$ pixels) into 12×8 boxes (each 1024×1024 pixels). Each star has a series of measurements at different locations on the mosaic. If the measurements are uncorrected, a given star will have a

large scatter, because measurements near the center of the mosaic are too bright, while those near the corners are too faint. Using an iterative process, we determined corrections for each of the 12×8 mosaic grid positions that minimized the scatter per star, and at the same time we determined best-fit magnitudes for each star, based on the collection of adjusted measurements. Stars were excluded from the analysis if they had intrinsic errors greater than 0.04 mag. The resulting 12×8 grid is converted to a full-resolution mosaic image by interpolating between the pixels.

Figure 8 shows the outcome of this analysis for the *R* filter. The top left plot shows the uncorrected stellar residuals as a function of the *X* mosaic coordinate, while the bottom left plot shows the stellar residuals after the grid of corrections is applied. It is clear that the measured corrections remove the large systematic trend. Compare the shape of the residuals in the bottom plot to those observed in the standard star data set (Fig. 4). The right-hand panel shows histograms of the residuals from the corrected data. The larger histogram shows the distribution of the residuals for all stars, while the smaller histogram shows the distribution for residuals of stars with magnitudes under 16, for which the Poisson errors should be less than 1%. The smooth curve overlapping this histogram is a Gaussian with $\sigma = 0.01$ mag. The formal scatter of these magnitude-selected

TABLE 2
STATISTICS OF ELIXIR IMAGE PROCESSING (FALL 2002)

Quantity	Number
CFH12K runs processed	26
Dark frames	74
Flat-field frames	146
Science frames processed	24,242
Frames distributed	5025
Sky pixels	2.4×10^{12}

residuals is 0.0086 mag. Clearly, the remaining systematic error per measurement is less than 0.01 mag. The results for the other CFH12K filters are similar to those for *R*. We have labeled this new correction “scatter-B.0” in the Elixir online documentation, and the corrected flat-field images are labeled “phot-flat-B.0.”

Data distributed for the CFHT QSO system since 2002 April have had the scatter-B.0 correction applied for both CFH12K and MegaPrime. We also document on the Elixir Web site these changes and provide recipes to convert flat-field images created with the scatter-A.0 correction to the scatter-B.0 correction. It should be noted that the two corrections, scatter-A.0 and scatter-B.0, are applied to the flat-field images differently: the scatter-A.0 is *subtracted* from the raw flat-field image, since it was constructed on the basis of the difference between flat-field frames, while the scatter-B.0 correction is *multiplied* by the raw flat-field frames, since it was constructed on the basis of stellar magnitude differences, which are flux ratios. The difference between these correction methods is due to the way the correction is measured, not the physical origin of the error that is being corrected; both corrections are compensating for the same errors in the flat-field. We can use a multiplicative correction here, because the ratio between the amplitude of the error and the amplitude of the flat field is extremely consistent over a wide range of twilight flat-field illumination levels. This consistency of these ratios is demonstrated by the consistency

of the (uncorrected) twilight flat-field images as the flux levels change over 2 orders of magnitude.

Figure 6 compares the *R*-band scatter-A.0 and scatter-B.0 corrections as gray-scale images. The full range of the gray-scale images is equivalent to a 1% correction to the flat-field image. The two patterns are generally similar, but there is somewhat more structure in the scatter-B.0 correction image.

6. SUMMARY AND CONCLUSIONS

The Elixir system has been in regular operation for QSO data since 2001 January, although several systems were introduced earlier. The first distribution processing was completed in 2001 September, at which point all stages of the reduction pipeline were functioning. Table 2 lists the Elixir processing statistics for CFH12K as of fall 2002. Since that time, we have learned a great deal about the operation of such a system, as well as about the CFH12K imager. In addition to running the Elixir system on all images obtained in QSO mode, all non-QSO images have also passed through the Elixir system, as well as all archived images since 1999 September. As a result, we have produced master detrend frames for all CFH12K runs since 1999 September, as well as an analysis of the standard star zero points. These images are applicable to the archived data, are available on CFHT’s Web site, and are being made available for distribution by the Canadian Astronomy Data Centre, which is also responsible for archiving raw CFHT images and will be the distribution center for the CFHT Legacy Survey with MegaCam.

We acknowledge the efforts of the entire CFHT New Observing Program: QSO, TCS, NEO, and DADS, and all of the staff at CFHT. We thank the members of the CFHT Scientific Advisory Committee and Board of Directors for helpful comments and guidance in the development of this project. We also thank the staff at both Terapix and CADC for their discussions and suggestions.

REFERENCES

- Bertin, E., & Arnouts, S. 1996, *A&A*, 117, 393
 Cuillandre, J.-C., Luppino, G. A., Starr, B. M., & Isani, S. 2000, *Proc. SPIE*, 4008, 1010
 Cuillandre, J.-C., Magnier, E. A., Isani, S., Sabin, D., Knight, W., Kras, S., & Lai, K. 2002, *Proc. SPIE*, 4844, 501
 Fahlman, G., & Grundseth, B. 2001, *CFHT Bull.*, 43, 2
 Forveille, T. 2001a, *CFHT Bull.*, 43, 14
 ———. 2001b, *CFHT Bull.*, 43, 18
 Høg, E., et al. 2000, *A&A*, 355, L27
 Landolt, A. U. 1992, *AJ*, 104, 340
 ———. 2004, *AJ*, submitted
 Magnier, E. A., & Cuillandre, J.-C. 2004, *PASP*, submitted
 Manfroid, J., Selman, F., & Jones, H. 2001, *Messenger*, 104, 16
 Martin, P. 2001, *CFHT Bull.*, 43, 2
 Martin, P., Magnier, E. A., Cuillandre, J.-C., Isani, S., & Fahlman, G. 2000, *CFHT Bull.*, 42, 2
 Mateo, M. L., & Schechter, P. L. 1989, in *Proc. First ESO/ST—ECF Data Analysis Workshop*, ed. P. J. Grosbol, F. Murtagh, & R. H. Wamels (Garching: ESO), 213
 Starr, B. M., Luppino, G. A., Cuillandre, J.-C., & Isani, S. 2000, *Proc. SPIE*, 4008, 1022
 Swade, D. A., Hopkins, E., & Swam, M. S. 2001, *ASP Conf. Proc.* 238, *Astronomical Data Analysis Software and Systems X*, ed. F. R. Harnden Jr., F. A. Primini, & H. E. Payne (San Francisco: ASP), 295
 Veillet, C. 2000, *CFHT Bull.*, 42, 30
 Withington, K., & Grundseth, B. 2001, *CFHT Bull.*, 43, 5