Duty Cycle Metrics Pilot Program

Final Report

12 September 2002

R. W. Goodrich

TABLE OF CONTENTS

1. Measuring observing efficiency

How efficiently do we collect science data?
How do we measure this efficiency?
How can we improve it?

These are some of the questions we must ask and answer if we want to make optimum use of the 10 m Keck telescopes and their related facilities. As an initial attack on these questions we have developed a pilot program with limited scope and time scale, the "Duty Cycle Metrics Pilot Program." This is the final report for that pilot program, being presented at its end.

2. Types of efficiency

We broadly classify the efficiency into three classes:

- *Duty cycle*
  What fraction of time is spent with the science shutter open? By minimizing the other tasks of the telescope and instrument we hope to maximize the duty cycle of the science detector. For example, we might speed up the telescope slew motors so that we can move from star to star more quickly. We might buy or construct faster readout electronics to minimize the amount of time it takes to get data from the instrument to the computer. We might create more efficient observing techniques, to allow observers to spend less time carefully centering their targets in a spectrograph's slit.

- *Throughput*
  How many of the photons raining down on the primary mirrors make it through to the science detector? This includes the effects of the telescope mirror coatings (primary, secondary, and when relevant, tertiary) as well as the reflective and absorptive losses in the instruments. We also include under this umbrella the science detector characteristics, which include not only the quantum efficiency of the detector but also, e.g., the readout noise.

- *Image quality*
  How efficiently does the telescope / instrument combination provide photons to the science detector. Image quality can often be tied to throughput, such as when good image quality at the focal plane of the telescope allows more light to pass through the slit of a spectrograph.

> But there are more subtle influences as well. In extreme cases poor image quality can prevent some science from being done at all!

The duty cycle metrics pilot program deals only with the first of these three classes.

3. The Pilot Program

The duty cycle metrics pilot program is the initial step towards an Observatory-wide observing efficiency program. The idea is to define and measure the fraction of time spent each night in various observing activities, such as slewing to a new target, centering it, and taking science data. The goal is to maximize the quality of the science data while minimizing the time and effort to obtain it.

3.1 Major goals of the pilot program

- Define the observing activities that will be used to delineate different phases of the night. These are given in Appendix B.
- Identify the "events" that distinguish the transition between different activities. These are given in Appendix C.
- Create a baseline of data for future comparison.
- Prototype data-gathering techniques, analysis, and visualization tools.
- Identify deficiencies in our current information sources.

3.2 Major goals of the full duty-cycle metrics program

- Reliably and sensibly track the duty cycle of observing activities.
- Identify significant sources of lost efficiency.
- Organize and advise efforts to ameliorate these efficiency losses by improving software, hardware, and electronics.

4. Pilot Program Details

An important part of any duty-cycle metrics program is the definition of the observing activities. To identify such activities we generally must rely on indirect sources of information. We *could* require the OA and / or observer to use the equivalent of a software "chess clock," with one clock for each activity. However, this would undoubtedly add to the observing overhead, and would be prone to human error.

Instead, we rely on sources such as time-stamped log files and image headers to gather information on various events that represent the start or stop of a particular observing activity.

4.1 Study log files for "events"

Events are considered "atomic" pieces of information. They are (generally) gleaned from various log files and indicate things like "guiding started / stopped / paused / resumed," "telescope slew started / stopped", etc. Almost all types of events have both a "start" event and a "stop" event or the equivalent.

An important exception is science exposures. The events for these are generally taken from the image headers, including information that determines whether they are calibration or science frames, imaging or spectroscopic frames.

A full list of events tracked and their sources is given in Appendix C.

4.2 Convert "events" to "activities"

The events taken from various log files are then ordered in time. This time-ordered sequence is analyzed to identify the start and stop times of various observing "activities." In contrast to events, activities are the familiar observing tasks like acquiring a new target, centering it on a spectrographic slit or coronographic spot, and taking data. Also, unlike events, observing activities are mutually-exclusive. The conversion from events to activities is context-sensitive; the interpretation of a given event may depend on the events that came both before and after it. A full list of the observing activities and their definitions is given in Appendix B.

It was argued early on that the software to extract events from the log files should be written in such a way that it would allow for ex post facto analysis of log files. This has two great advantages:

- It allows immediate access to much of the metrics information from the previous year or more, providing an instant baseline that can be used for comparison.
- It allows changes to the software "rule set" to be applied to existing log files, rather than having to restart the entire metrics data set from scratch. To illustrate this, let's assume that you had made a mistake and were triggering the "start slew" event off the wrong line in the telescope log file. If you had no ability to go back to existing log files,

once you fixed the problem with the “start slew” event you would have to start your metric baseline from scratch again.

Note that while images are not kept on disk for more than 14 days, their headers, from which the metrics information is extracted, are filed on disk for several years, and hence are readily accessible.

The atomic “events” are written to flat ASCII files for ease of access by different and varied tools.

During this process and the next (conversion to observing activities) we make note of desired or required events that are not currently logged adequately. These inadequacies must be addressed by the full metrics program.

4.3 Calculate duty cycles for the observing activities

The fraction of the night spent in various observing activities is calculated. Note that “night” is defined as the time between 12° twilights.

The conversion can either be done “on the fly” by reading the relevant ASCII event files, or the conversion can be done at the time the event files are created and stored in a database for more rapid retrieval.

4.4 Analyze duty cycle distributions for different instruments/modes

Over a long baseline (at least six months), analyze the average and extreme values of the duty cycle values for different instruments and their different observing modes. When problem areas are defined and addressed, this baseline provides a measure by which any efficiency improvements can be judged.

It is important to differentiate between different observing modes on a given instrument. A good example is LRIS, the Low-Resolution Imaging Spectrograph. We can consider three different observing modes with LRIS: slitmask spectroscopy, long-slit spectroscopy, and imaging. For slitmask spectroscopy, we might expect relatively large overheads in “fine acquisition,” while the objects are precisely centered in the slitlets. Scientifically this is compensated by the large number of spectra that can be taken in one exposure.

For imaging, especially with broadband filters (UBVRI), we might expect that the science involves relatively short exposure times per frame. This may create a higher overhead due to the readout time of a frame.

With some instruments, such as NIRC, LWS, and NIRC-2, centering a target on a spectroscopic slit involves taking images with the science detector while adjusting the telescope position. With NIRSPEC, however, one generally centers the target using the slit-viewing camera, which has a relatively shorter readout overhead and frame rate. The relative ease with which centering can be done impacts the observing efficiency for these infrared instruments and may influence future designs or retrofitted changes to the existing instruments.

Note that the limited time scale of the pilot program has obviated the usefulness of these long time scale comparisons *for the pilot program*, but the ability to look back in time and get metric information from historic log files makes the time period of the pilot program useful for the full metrics program.

4.5 Create visual feedback for each stage

It is important for the human interpretation of these results to provide straightforward, readily-interpreted visual feedback, such as charts and graphs. These visualization tools are used to confirm that the event definitions make sense, that they are properly interpreted and converted into appropriate observing activities, and that the observing activities are properly tabulated.

Implementation of the visualization tools has included the used of Web-based graphical representations of events and activities.

5. Pilot Program Results

5.1 Events

Most events are logged in one of several places:

- The autoguider log file, autoguider.syslog.
- One of the telescope log files, envCmd.arM.
- Various instrument log files.
- Instrument images FITS headers.
- OA-entered fault times in the Remedy database.

Some events are not logged adequately:

- Weather time lost. (The total time is logged, but not the time period.)
- Instrument reconfiguration start and stop times.

These inadequacies are addressed in the Recommendations section of this document.

Note that faults can occur in any number of systems, and may or may not result in logged information. Hence we rely on the OAs (Observing Assistants) to identify the time period affected by a given fault.

5.2 Conversion of events to activities

Sometimes this conversion is ambiguous. For example, is the target being observed as a science target or a calibration target? Are images taken by the science camera being taken to aid in fine alignment, or for scientific purposes?

We have chosen to be relatively conservative in our interpretation, adhering to the statement, “Assume it is science unless you can show otherwise.”

An exception to this is that images of a target taken before a spectrum of that same target are assumed to be for fine alignment purposes, not for science. This is because it is a standard technique used for fine alignment in five of the seven instruments covered. LRIS’s slitmask observing mode is one of the most common of these circumstances, in which images are taken to align the masks to the sky before taking the slitmask spectra that are the real scientific data product. NIRC, NIRC-2, and LWS all use images with their science detector to align the target on the slit. NIRSPEC uses images taken with the slit-viewing camera (SCAM) to center science targets on the spectrograph’s slit. (SCAM can also be used for scientific imaging at times.)

In order to implement a rule to cover these common observing scenarios, it is essential that the rules for converting events to activities include the ability to “look back” and see what has occurred before a given event. Thinking in terms of the time-ordered event sequence, if a slew and coarse acquisition is followed by images taken by the science instrument, these are initially classified as science data (“open shutter time”). However, if a subsequent spectrum of this same target is taken, the conversion routine must “look back” to the earlier images and reclassify them as fine acquisition. This context-sensitive rule set is essential to the success of the classification.

The requirement that the observing activities be mutually exclusive implies an order of preference for the different activities. For example, if the telescope is pointed at a target, and data is being taken, this would normally be classified as science data. However, if this is during a time period classified as “bad weather,” the weather classification takes precedence and the period is classified as weather rather than science. Such a scenario can occur if the observers are taking calibration data on the sky while waiting for weather to clear.

Similarly, if a fault occurs during a period of good weather, this classification would take precedence over all other classifications. In order to diagnosis or fix a fault, the telescope may have to be slewed to a new position, or the observer or support person may have to record frames with the instrument that would otherwise look like science data to the event-activity conversion algorithm.

5.3 Visualization tools

Important to the overall metrics effort are tools to aid in analyzing and interpreting the information. In the pilot program we have developed a number of useful visualization tools:

5.3.1 Nightly event time line

This displays events from various log files on a labeled time line, as extracted by the event software. The tool allows for the selection of different time base lines within a given night, allowing the user to concentrate on certain small parts of the night. The tool provides a graphical representation of the event sequence and context in order to aid in interpreting the event-activity conversion.

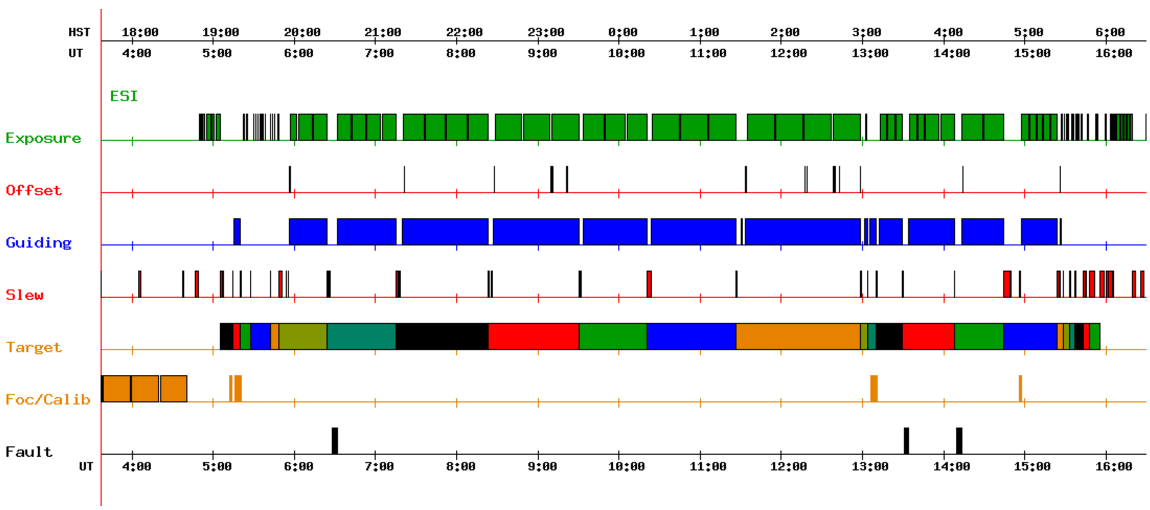

Figure 1. An example of a nightly event time line, in this case for ESI.

5.3.2 Nightly activity time line

Once events are converted to observing activities, this tool displays the activities in a very analogous way to the event time line. Again, the user can specify a narrower time range in order to concentrate on one part of the night. This tool provides a graphical representation of the results of the event-activity conversion.

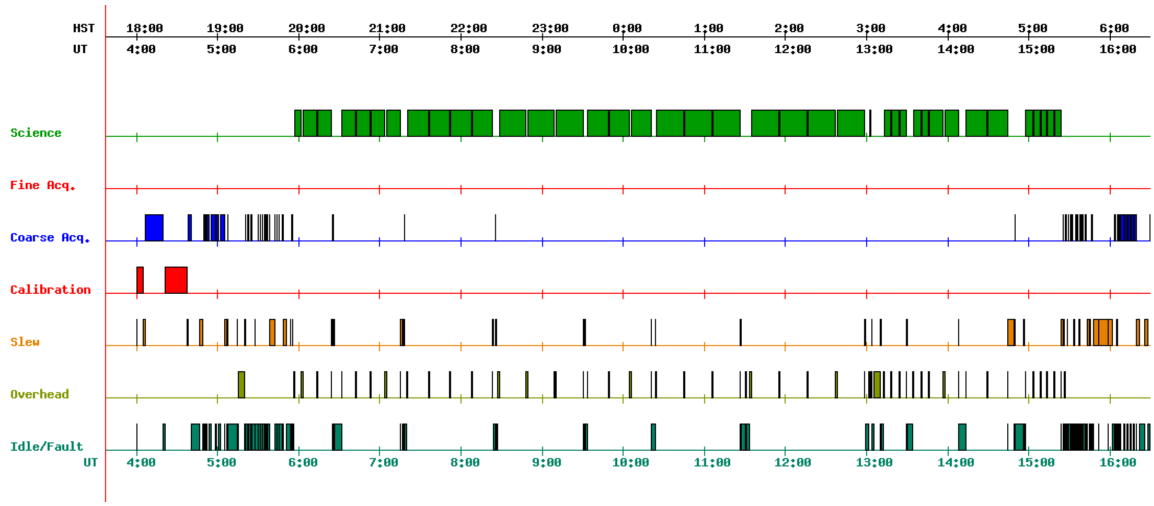

Figure 2. An example of an activity time line. Note that the rule set for conversion from events to activities was not finalized when this chart was created, so it is not 100% accurate, but provided as an example only.

5.3.3 Duty cycle pie chart

After calculation of the fraction of the night spent in the various observing activities, a pie chart is created. This provides a quick visual summary of the total time spent in each activity, with the eventual goal of increasing the size of the slice representing science, “open shutter” time, and decreasing the sizes of the slices representing other observing overhead.

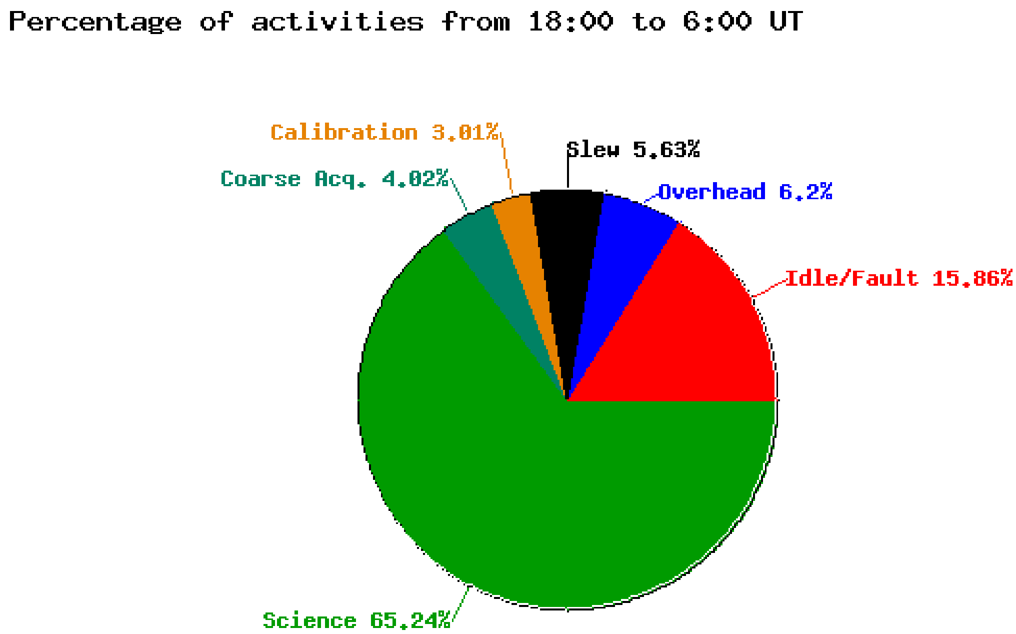

Figure 3. A pie chart summarizing the fraction of time spent in various observing activities. This was derived from Fig. 2, and hence is also not 100% accurate.

5.3.4 Other, unimplemented tools

Not implemented in the pilot program, mostly because of its limited scope and time scale, are other recommended visualization products, such as: pie charts showing monthly averages, pie charts showing averages by instrument and observing mode, and similar graphs showing not just the averages, but also the night-by-night minimum and maximum for a given month or instrument. It is recommended that these tools be implemented in the full metrics program.

5.4 Notes

Some important limitations to the pilot program should be mentioned:

5.4.1 No "real-time" system

The pilot program was *not* developed as a "real-time" system. In a real-time system the events would be recorded as they occur, the observing activity defined from the event sequence (looking backwards only in this case), and the visualization tools could be updated in real time.

Instead, the path chosen was one of waiting until the night is over, making a single pass to get the event information for the night, converting the event sequence to observing activities, and presenting the information from there.

A real-time system was not deemed necessary for the goals of the full metrics programs Its development also would have detracted from the development of other tools during the pilot program. While a real-time system would have been flashy and of interest to some of the OAs and observers, the true importance of the metrics program lies in the accumulation and analysis of statistically significant amounts of data, not in the analysis of a single night's observing pattern. A real-time system of limited utility will also just add one more piece of software that must be maintained.

5.4.2 Adaptive optics

In the pilot program, adaptive optics (AO) was not investigated. This is significant because the AO system "usurps" some telescope (DCS) functions like guiding. During an AO night the usual DCS autoguider log file will not show any "guiding" activities. The AO equivalent would be taken from events in the AO log files, and would consists of the start and stop times of locking the "tip-tilt" loop.

Despite the length of time that AO has been in regular scientific use on Keck II, the observing tools and techniques are still evolving significantly. Thye will continue to do so in the future. This may hamper development of algorithms identifying events in the log files or converting events to observing activities.

5.4.3 DEIMOS

The newly-installed DEIMOS instrument on Keck II is still undergoing commissioning. It was installed after the start of the pilot program, and was not included in the boundary conditions of that program. It is, however, very similar to ESI, as well as to LRIS and HIRES, so incorporation of DEIMOS information in the full metrics program should not pose a major problem.

During the ongoing commissioning and the initial science runs, DEIMOS observing techniques are still being refined. This creates something of a "moving target" for the metrics program. How this affects the metrics software remains to be seen.

6. Pilot Program Recommendations

The main product of the duty-cycle metrics program is a list of recommendations on how to implement the full duty-cycle metrics program. As a secondary benefit, some of the software prototyped during the pilot program may be useful in the full program.

The recommendations of the pilot program are:

- ✍ Retain the general structure of the pilot program:
  — extract “atomic” events from log files or image headers,
  — convert events to a limited number of observing “activities,”
  — tabulate nightly fractions of time spent in each activity.

- ✍ Refine the rule set for the event-activity conversion.
  Currently the rule set has not been developed sufficiently enough to provide accurate, robust conversion.

- ✍ Add new OA tools:
  — “weather” ticket type to log the start and stop times of periods lost to weather,
  — clear OA instructions to modify start time / time lost values for faults to correspond to actual time period affected. For example, if a 60-minute exposure is lost due to a fault, the start time of that fault should be logged as the beginning of that exposure, not the time that the fault actually occurred or was discovered. If no useful data was obtained on that target because of the fault, the start time would be advanced to the start of the slew to that target.

- ✍ Relieve the OA from some tasks.
  Currently the OA has to manually sum up some types of tickets, such as the “Time lost” fault tickets. This should be automated regardless of the metrics program, but in particular it will lessen the burden of the OA responsibilities added by the metrics program requirements.

- ✍ Investigate further a “real-time” system. This will be done at a lower priority, and as a precondition, the failure of any real-time system developed will not interfere with the normal collection and interpretation of the metrics data.

- ✍ Revise various pieces of instrument software to better log:
  — motor move start / stop times or other reconfiguration information,
  — overhead estimates related to each exposure.

- ✍ Maintain current visualization tools and develop others (see section 5.3):
  — nightly even time line (developed in pilot program),
  — nightly activity time line (developed),
  — single-night activity pie chart (developed),

— monthly activity pie chart,
— instrument / observing mode pie chart,
— monthly min / average / max chart,
— instrument / observing mode min / average / max chart.

✍ Create a duty-cycle “change log” to indicate when changes are made and what those changes are. Generally these would be changes made specifically to enhance the nightly scientific output, although other changes such as a decrease in telescope slew speed because of some mechanical problem in the drives should also be entered.

✍ Provide quarterly reports to the Keck SSC (Science Steering Committee) and CARA Board on progress in improving the observing duty cycle. Allow external access to at least some of the data.

✍ A man-power estimate is provided in Appendix A.

✍ Incorporate AO and DEIMOS.

✍ Based on regular analysis of the duty cycle calculations by the Keck metrics team, as well as the operations team (and to a lesser extent the software, electronics, and mechanical teams) , recommend courses of action to improve the scientific duty cycle, and prioritize the recommendations.

The “metrics core team” will meet every two months to study the effects of any previous changes in observing procedures, hardware, electronics, or software. They will also make recommendations for further improvements. These recommendations will be passed on to one or more of various teams within the Observatory, such as the support astronomers, electronics group, software group, and mechanical group. These groups, and the Observatory Director and Assistant Directors can, of course, reprioritize the suggestions, suggest alternate or additional improvements, and integrate the work into the groups’ own work plans.

✍ Maintain the present metrics Web site with information ranging from details on the implementation of various pieces of software, through punch lists of tasks to complete, to results of the metrics analysis.

Appendix A—Effort Estimates for the Full Duty-Cycle Metrics Program

Task	Details	FTE	Personnel
Data collection	Perl scripts	4 mo.	MT/RG
	adding new log entries	4 mo.	MT/EC/AH
	Remedy modifications	2 wk.	JS
	cron job maintenance	1 mo.	MT
Data processing	events/activities	2 mo.	MT
Data presentation	Web pages	2 mo.	SK
Reports	SSC/Board reports and Web reports	1 mo.	RG/SK/AC
Interpretation/ recommendations		1 mo.	RG/SK/AC

Personnel:

AC	Al Conrad
AH	Allan Honey
EC	Liz Chock
JS	Julia Simmons
MT	Myrna Tsubota
SK	Shui Kwok
RG	Bob Goodrich

Appendix B—Observing Activities

- Weather
  Observing was essentially on hold because of bad weather.
- Fault
  Some problem was preventing observing. This could be a telescope problem, an instrument problem, a dome problem, etc.
- Slew
  The initial telescope move to a new target.
- Instrument reconfiguration
  Motors are moving in order to change settings of the instrument.
- Detector overhead
  The time it takes to erase or readout a CCD, or reset an infrared detector.
- Calibration
  Exposures such as flat fields that will be used for calibration, but not directly for science.
- Science
  The “holy grail,” this is the “open shutter” time that we wish to maximize.
- Coarse acquisition
  The initial identification and rough positioning of a new target.
- Fine acquisition
  Small movements to accurately center a target, often in a slit. Unlike coarse acquisition, guiding is turned on during fine acquisition.

Appendix C—Events and their Sources

Event	Source
Slew start / stop	envCmd.arM (DCS log file)
Guiding start / stop	autoguider.syslog (DCS log file)
Offset start / stop	autoguider.syslog (DCS log file)
Target change	envCmd.arM (DCS log file)
Instrument change	envCmd.arM (DCS log file)
Guide camera change	autoguider.syslog (DCS log file)
Science exposure start / stop	image FITS headers
Calibration start / stop	image FITS headers
Focusing start / stop	focus log
Fault start / stop	OA entries in Remedy database