

Semi-weekly monitoring of the performance and attitude of *Kepler* using a sparse set of targets

Hema Chandrasekaran^{a,b}, Jon M. Jenkins^o, Jie Li^a, Forrest R. Girouard^c, Joseph D. Twicken^a, Douglas A. Caldwell^a, Christopher Allen^c, Stephen T. Bryson^d, Todd C. Klaus^c, Miles T. Cote^d, Brett A. Stroozas^e, Jennifer R. Hall^c, Khadeejah Ibrahim^c

^aSETI Institute/NASA Ames Research Center, M/S 244-30, Moffett Field, CA USA 94305

^bLawrence Livermore National Laboratory, P.O. Box 808, L-478, Livermore, CA USA 94551

^cOrbital Sciences Corporation/NASA Ames Research Center, M/S 244-30, Moffett Field, CA USA 94305

^dNASA Ames Research Center, M/S 244-30, Moffett Field, CA USA 94305

^eStroozas FlightOps, 420 Augustus Ct., Walnut Creek, CA USA 94598

ABSTRACT

The *Kepler* spacecraft is in a heliocentric Earth-trailing orbit, continuously observing $\sim 160,000$ select stars over ~ 115 square degrees of sky using its photometer containing 42 highly sensitive CCDs. The science data from these stars, consisting of ~ 6 million pixels at 29.4-minute intervals, is downlinked only every ~ 30 days. Additional low-rate X-band communications contacts are conducted with the spacecraft twice a week to downlink a small subset of the science data. This paper describes how we assess and monitor the performance of the photometer and the pointing stability of the spacecraft using such a sparse data set.

Keywords: *Kepler Mission*, attitude reconstruction, exoplanet, transit, tracking and trending, performance monitoring

1. INTRODUCTION

The *Kepler Mission* continuously observes $\sim 160,000$ target stars in *Kepler's* 115-square-degree field of view (FOV), seeking to discover Earth-like planets transiting Sun-like stars by detecting photometric signatures of transits.^{1,2} The photometer's 42 CCDs are read out every 6.52 s and co-added for 29.4 minutes*, after which the pixels of interest for each star are compressed and stored on board the spacecraft's Solid State Recorder (SSR). The data are downlinked at monthly contacts with the Deep Space Network (DSN). *Kepler's* primary mission will last 3.5 years to permit detection of at least three transits of Earth-size planets orbiting Sun-like stars in the habitable zone, that range of distances at which liquid water could pool on the surface. Only a small fraction ($\sim 5\%$) of the 96×10^6 pixels in each frame are stored, with a budget of 32 pixels per star, on average. In addition to the target star pixels, 4,500 background pixels are collected per CCD readout channel along with 1100 masked smear, 1100 virtual smear, and 1070 trailing black measurements.^{3,4} The masked smear and virtual smear allow us to estimate and remove the vertical smear trails due to the lack of a shutter, while the trailing black measurements allow us to estimate and correct for readout-dependent changes in the bias voltages in the analog signal processing chains. The total amount of data stored onboard the SSR is ~ 174 MB per day, and the SSR can store ~ 66 days of data. The science data are normally downlinked at Ka-band through the high-gain antenna (HGA), which is fixed to the spacecraft bus. Thus, *Kepler* must stop collecting science data at the science attitude and slew to point the HGA toward Earth to downlink science data.

A small amount of science data consisting of no more than 96,000 Long Cadence (LC) pixels are collected once per day and stored in non-volatile memory for downlink at the twice-weekly X-band contacts with the DSN to provide

^oFurther author information: Send correspondence to J.M.J.: E-mail: Jon.Jenkins@nasa.gov

©2010 Society of Photo-Optical Instrumentation Engineers. One print or electronic copy may be made for personal use only. Systematic reproduction and distribution, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper are prohibited.

*Each 29.4-min data integration interval is called a Long Cadence (LC); the data is called LC data to distinguish it from data collected at \sim one-minute intervals, called Short Cadence data.

frequent updates on the status and health of the photometer and to assess the validity of the spacecraft’s attitude for science operations. Up to nine reference pixel files can be stored on board, so that in the event of a missed X-band contact, all the reference pixel files can be downlinked during the subsequent DSN pass. These “reference pixel” data are used by the Science Operations Center⁵ (SOC) pipeline component Photometer Data Quality (PDQ) to produce metrics related to the health and performance of the photometer, such as measurements related to the bias level, background flux level, dark current, photometer throughput, focus, and pointing stability and accuracy. Similar metrics are monitored by the SOC software module Photometer Performance Assessment (PPA) which can only be run after new science data have been downlinked to the DSN at Ka-band nominally once per month.⁶ Table 1 contains a list of the principal metrics calculated by PDQ and their purpose. PDQ tracks and trends these metrics throughout each 93-day quarterly segment to identify sharp discontinuities in their behavior, or long-term drifts that cross pre-determined hard bounds. In essence, PDQ plays the role of the “canary in the coal mine” to allow us to identify either abrupt changes in behavior that may signal hardware failures or issues, or to spot gradual degradation in the performance of the spacecraft or photometer that might otherwise not be identified in the tracking and trending of the engineering data by the Mission Operations Center (MOC). The SOC Operations staff run PDQ when a new set of reference files arrive at the SOC and notify the Science Office (SO) of the arrival and processing status of the files.⁷ When the processing is complete, the report files and all the task files, including input files, intermediate data products, and output files, are copied over to a server accessible by SOC and SO staff for analysis. The PDQ reports contain a plethora of information and graphics on the status of the photometer and are broken into a summary report for the entire focal plane and 84 individual reports for each CCD readout channel. Figure 1A shows the “dashboard” plot for one PDQ run during Q3 observations.

Table 1. Performance Metrics Calculated by PDQ

Metric Name	Units	Purpose
Dynamic Range	ADU	Simple check for “aliveness” of each CCD readout channel
Black Level	ADU	Monitors drift in CCD readout channel bias voltages
Smear Level	e ⁻	Monitors the average shutterless readout smear level
Dark Current	e ⁻ s ⁻¹	Monitors drift in dark current
Brightness	n/a	Monitors change in photometer throughput
Residual Centroid Rows & Columns	pixels	Monitors drift in pointing
Encircled Energy	pixels	Monitors focus and high-frequency pointing jitter
Plate Scale	arcsec	Monitors drift in photometer focus
Photometer Attitude (ra, dec, roll)	degrees	Reconstructed attitude
Maximum Attitude Residual	pixels	Monitors pointing discrepancies

Module 3 of the Focal Plane Array (FPA) died abruptly on January 9 of 2010. Its failure was identified immediately when PDQ was run on the four new reference pixel files brought down on the next X-band contact. The availability of science data from the two affected CCDs assisted in the diagnosis of the likely mode of failure in the immediate aftermath and also heralded the continued healthy status of the remaining 40 CCDs on 20 modules in the FPA.

One of the most important tasks that PDQ supports is the analysis of the spacecraft’s pointing attitude following a break in science operations to perform a quarterly roll maneuver to reorient the sunshade and solar arrays, to downlink data from the spacecraft, or due to an unplanned safe mode event. The pointing precision required of the Attitude Determination and Control System (ADCS) is 0.009” at 3σ on timescales of 15 min and longer,² whereas the absolute pointing knowledge required is at the 12” level (3σ). The pixel budget for the target stars is rather tight, permitting only a one-pixel “halo” around the pixels that are needed for the photometric measurements. PDQ is required to reconstruct the spacecraft’s actual attitude and compare that to the nominal one, furnishing a pointing “tweak” capable of bringing the pointing to better than 0.1 pixel or 0.4” of the desired attitude across the focal plane.




Figure 1. Left: Summary dashboard plot for bounds crossings of all metrics across the focal plane is shown, giving the user a complete picture of the general health and performance of the instrument and spacecraft from a science perspective. Right: The information is arranged as per this key. The color scheme is as follows: Green - no bound crossings over the cadences from the start of the latest contact; Yellow - crossed adaptive bounds at least once over the cadences; Red - crossed fixed bound at least once over the cadences; Cyan - no data available. Metrics related to the pointing are presented in the 1×4 bar in the upper left hand side of the dashboard. These include the boresight right ascension, declination, photometer roll angle, and the maximum attitude residual (see Sec. 4). The other metrics are computed for each CCD readout channel as listed in Table 1). This plot was generated after Module 3 died on January 9, 2010, so the metrics for all four CCD readout channels on Module 3 are colored cyan, to indicate that no valid data is available.




Figure 2. This graphic shows the locations of the reference pixels used to monitor the health and performance of channel 2.1 during Q2. A: An overlay of the pixel locations on a Full Frame Image (FFI) taken just before a monthly downlink. There were three stars collected as reference pixel targets, along with their associated collateral data, five background pixels, masked (rows 7–11) and virtual smear (rows 1051–1055) for each column, and trailing black measurements (columns 1127–1131) for each row containing a star or background pixel. Pixels are also collected for each row of smear in the trailing black region to correct for bias. B: A graphic from the PDQ report for channel 2.1 showing the uncalibrated target star pixels on the top row and the calibrated target star pixels on the bottom row for the three reference pixel target stars on this channel.




Figure 3. The processing steps in PDQ applied to reference pixel data to obtain the various desired metrics associated with each CCD readout channel. The calibration steps mirror those used for normal science data.

2. PROCESSING REFERENCE PIXELS

Reference pixel files are downlinked from the *Kepler* spacecraft twice weekly through the DSN to the MOC at the Laboratory for Atmospheric and Space Physics in Boulder Colorado, which forwards the files directly to the SOC at NASA Ames Research Center. Each reference pixel file contains one 29.4-min sample for no more than 96,000 pixels in the focal plane. Typically each CCD channel has 3–5 stellar targets chosen so that they are well distributed across the channel, along with several background pixels, plus collateral data consisting of virtual and masked smear pixels for each column and trailing black pixels for each row that contains a target or background pixel. Unlike normal science target stars, collateral data must be specified through the use of stellar target table definitions.⁸ Figure 2 shows the locations of the pixels used for module 2 output 1 (the first channel) during Quarter 3 (Q3)[†]. The calibration steps in PDQ follow those for use in calibrating nominal long cadence science data,⁹ as shown in Figure 3. Reference pixel data is distinct from normal science data both in terms of how the collateral data is specified and in terms of the required output products. The Science Processing Pipeline used to process normal science data cannot handle reference pixel data, necessitating special-purpose software to meet unique data format and output product requirements. PDQ is essentially the “SOC in a box” in that it calibrates the pixels in the reference pixel file, extracts photometry and astrometry from the stellar images, reconstructs the pointing, and tracks and trends a number of metrics related to the performance of the photometer. There are models¹⁰ for all the photometer characteristics relevant for calibrating and correcting *Kepler* science data, and these are described in detail in Ref. 3.

The reference pixel data are processed one CCD channel at a time. The first metric calculated from the raw pixels in each channel is the dynamic range metric. This metric is simply the difference between the maximum pixel value and the minimum pixel value for each CCD channel. As the spacecraft drifts farther from Earth each year, the bandwidth available at X-band drops, requiring us to reduce the amount of reference pixel data downlinked. Although our first choice is to reduce the number of reference pixel files downlinked at each contact, eventually we will be forced to reduce the total

[†]Q3 began on September 18, 2009 at 17:19 UTC and ended on December 16, 2009 at 23:55 UTC.

number of pixels in each reference pixel file significantly and will not be able to collect star data on all the CCD channels. The Target Aperture Definition (TAD) pipeline component allows us to request a few pixels on each CCD channel from which the dynamic range metric can be calculated.⁸ As was unfortunately proven by the loss of module 3 earlier this year, the dynamic range metric is a good indicator of the basic health and status of the focal plane electronics. This metric is unique in that PPA does not calculate it for normal science data.

The first correction to the reference pixels is the removal of the 2-D black level or bias level. This is a static 2-D table with a value for each pixel in a Full Frame Image (FFI) per read. Fixed pattern noise is apparent in “dark” frames due primarily to electronic crosstalk from the Fine Guidance Sensor (FGS) clock driver signals.³ Since the *Kepler* electronics is a synchronous system, the FGS clock signals are in phase for any given science pixel, making this phenomenon static to first order. FGS clocking crosstalk is weakly dependent on the temperature of the electronics, and PDQ tracks and trends changes in the mean black level. Changes in the details of the 2-D black are tracked by the PPA on monthly LC data. The static 2-D black value for each pixel is multiplied by the number of readouts and subtracted from the raw pixel value. The trailing black pixels, which are samples obtained by over-clocking each row read out from each CCD readout channel by 20 extra pixels, are then used to estimate the residual black level or bias voltage. As this may be a function of the readout row number, a low-order polynomial is fitted robustly to the trailing black pixels as a function of row and then evaluated to form the 1-D black level estimate. The average of the robust polynomials fitted to the residual trailing black pixels on each frame is the *black level metric*.

The black-corrected pixels are then converted from analog-to-digital units (ADU) to units of photoelectrons (e^-) by multiplying by the channel-specific gain in $e^- \text{ ADU}^{-1}$. The second major calibration step is to apply the Local Detector Electronics (LDE) undershoot correction. LDE undershoot occurs due to the limited bandwidth of the CCD video output and a somewhat weaker than expected sample-and-hold circuit. This means that each pixel value digitized by the analog-to-digital (A/D) converter is really the weighted sum of the analog pixel value and previous ones. The undershoot is small, less than 1% for the first lag, and drops off rapidly thereafter. Still, it is a measurable effect in the images, particularly after a saturated pixel is read out. The remedy is to apply a digital filter to the data obtained in each CCD row by a realization of the inverse of the forward LDE undershoot filter[‡].

The third correction step is to estimate the smear from the shutterless readout and the dark current from the calibrated masked smear and virtual smear pixels. The former are physical CCD pixels that are prevented from receiving photons by an opaque mask positioned over the first 20 rows of each CCD. The latter are obtained by over-scanning the CCD by 26 rows following the readout of the physical CCD, which is 2200 columns by 1044 rows. Each pixel in a column picks up excess flux from stars and other flux sources above it as it is read into position during the previous readout (~ 0.5 s) for the 6.02 s exposure, and then picks up excess flux from stars and other sources below it during the subsequent readout. To first order, this excess flux is the same for each pixel in each column. The principal difference between the masked and virtual smear is that the latter do not exist during the exposure, and so collect only 0.5 s of dark current during the readout. Thus we can estimate the dark current for each channel separately from the smear level (in each column) by taking appropriate linear combinations of the masked and virtual smear measurements. The smear and dark current are then subtracted from each target and background pixel, and the dark current and median smear level are tracked and trended and reported as the *smear and dark level metrics* in the PDQ report.

Following the smear and dark level correction, the target and background pixels are corrected by the flat field model, which represents the sensitivity of each pixel in the focal plane relative to the mean pixel sensitivity. The background for the target star pixels is estimated from ~ 5 background pixels specifically chosen for this purpose together with dim pixels in the periphery of each stellar aperture mask. There are generally more good background pixels available from the masks since reference pixel masks have a “double” halo so that PDQ is robust against pointing acquisition errors. A robust average of the background pixels selected for each target star is subtracted from each pixel in a target star’s mask. The *background level metric* is taken to be the median of the background pixels on each channel.

The background-subtracted pixels in each target star’s mask are used to estimate the flux and the row and column centroid (photocenter) for each reference pixel file. The flux is calculated by adding up the pixel values in the “optimal aperture” identified for each target star by the TAD pipeline component.⁸ The centroid of each target is obtained by fitting a super-resolution representation of the point spread function convolved with the response of a single pixel, which is called a pixel response function (PRF).¹¹ The PRF-fitted astrometry is less biased and more robust than simple brightness-weighted

[‡]The LDE undershoot is monitored by PPA, but not by PDQ.

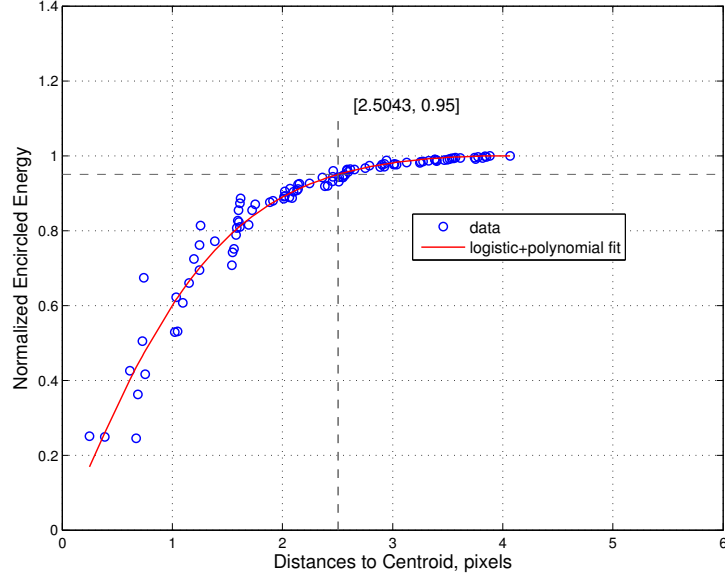


Figure 4. Encircled energy calculated for target stars on channel 2.1 during Q3 observations. The pixels for each target star are normalized by the total flux and then sorted by distance from the measured centroid, normalized by the maximum distance of all measurements. The pixel values for each star are summed cumulatively, so that the resulting values represent the fractional flux contained within a radius equal to the distance to each pixel. A hybrid logistic function and polynomial defined by Eq. 1 and constrained by the conditions in the text is fitted to the data points for all stars. The radius at which the fitted function passes through 0.95 is found and defines the encircled energy metric for this time step.

centroids. The flux measurements for each target star are normalized by the fraction of the flux expected in the optimal aperture and by the flux expected according to the target star’s *Kepler Input Catalog* magnitude, which is an estimate of the brightness of each star in the *Kepler* photometer’s bandpass.¹² The *brightness metric* is taken as the robust mean of the normalized fluxes for all targets on each CCD readout channel.

The final metric calculated on a CCD readout channel basis is the *encircled energy metric*, a measure of the focus. This metric is defined as the radius within which 95% of the flux is contained for the image of a star, and is calculated by sorting the pixels for each target by distance from its PRF-fitted centroid, then normalizing the sorted pixel values by the total star flux. The maximum distance from a centroid to any pixel included in this analysis is used to normalize all distance measurements to the interval $x \in [0, 1]$. The sorted pixel values are then summed cumulatively. An analytic curve is fitted as a function of the distance from the centroids to the normalized accumulated pixel values for all the targets on a CCD readout channel. The analytic curve is the combination of a logistic function and a cubic polynomial

$$f(x) = \frac{1 - \exp(-\beta x)}{1 + \exp(-\beta x)} + c_0 + c_1 x + c_2 x^2, \quad (1)$$

where $f(x)$ is constrained at the endpoints to meet the conditions $f(0) = 0$, $f(1) = 1$, and $f'(1) = 0$. The term β controls how quickly $f(x)$ rises from 0 at $x = 0$ to 1. The radius where the fitted analytical function attains a value of 0.95 is taken to be the *encircled energy metric*, as illustrated by Figure 4.

The metrics calculated by PDQ are tracked and trended to identify outliers and step-changes in behavior, as discussed in the next section.

3. TRACKING AND TRENDING

The PDQ metrics are trended using a simple moving average with an exponential forgetting factor, as detailed in Ref. 6, and a slight variation of this algorithm is used to estimate the variance of each metric. This moving average allows PDQ to “learn” the natural variance and value of each metric and adapt as the thermal environment of *Kepler* changes over time.

The forgetting factor can be adjusted to balance the memory of the system and the amount of smoothing supplied by the algorithm. “Adaptive” bounds are initially set quite broad for each metric, but as the variance estimate improves over time, these adaptive bounds shrink to provide more sensitivity. If an adaptive bound is broken, a “yellow” alert is logged. Yellow alerts are fairly common, especially after an operational activity such as a monthly downlink, due to the disturbance of the thermal state of the spacecraft and instrument. Fixed bounds are provided through operations-controlled parameters. If any metric breaks the fixed bounds, a “red” alert is logged. Some red alerts occur due to known features of the electronics used to drive and read out the CCDs. The SO monitors the alerts for unexpected and unexplained ones, which are reported to the *Kepler* Project Office for action.

PDQ also forecasts whether any fixed or dynamic bounds are likely to be broken in the future by fitting the last several days of each metric with a line and projecting the line forward in time to see how long it takes before a bound would be broken. If the projected time to break a bound is less than seven days (the “horizon time” – a configurable parameter), then an alert is logged and recorded in the PDQ summary report.

4. GENERATING AN ATTITUDE ADJUSTMENT

The pointing precision required of the Attitude Determination and Control System (ADCS) is $0.009''$ at 3σ on timescales of 15 min and longer. The absolute pointing knowledge, however, was initially only at the $12''$ level. Thus, analysis of some science data was necessary to furnish small attitude “tweaks” to control the pointing error at the initiation of science data collection to less than 0.1 pixels or $0.4''$. Since the science data are normally downlinked through the HGA, it was necessary to develop a way to downlink a sufficient amount of science data through the omni-directional, X-band antennas to permit an attitude analysis while the spacecraft is locked onto the FOV in fine point mode.

The centroids from all the available stars on all the CCD channels are aggregated together for each time step and used to reconstruct the pointing. Essentially, the right ascension (α), declination (δ), and roll angle of the spacecraft boresight (ϕ) are adjusted until the discrepancies between the measured and predicted locations of the target stars are minimized at the pixel level by a robust function-fitting algorithm, as detailed in Ref. 6. The difference between the attitude reconstruction performed in PDQ and that performed by PPA is that the latter uses so-called motion polynomials to represent the location of stars as measured by their centroids, to minimize the amount of data that is needed.

4.1 A Heuristic Pointing Model

Reconstructing the spacecraft’s pointing attitude from the measured centroids of a large number of target stars is straightforward given the spacecraft trajectory, the stars’ inertial coordinates, and the focal plane geometry model. This reconstructed pointing needs to be compared to the expected pointing at the time the data were taken.

Ten stars are tracked on each of the four fine guidance sensors that are located on the corners of the focal plane. These CCDs are read out at 10 Hz, and 8×8 “postage stamp” apertures containing the image of each of these FGS target stars are used to construct brightness-weighted centroids. Velocity aberration causes the center of the FOV to move in an ellipse $\sim 22''$ in radius as the spacecraft cruises in its orbit about the sun. Differential velocity aberration (DVA) causes the FOV to appear to contract as the spacecraft travels towards the FOV and to expand as the spacecraft travels away from the FOV over the course of *Kepler’s* orbit about the Sun. The FOV also appears to nod both clockwise and counter-clockwise slightly during quarters when the velocity aberration attains an extremum, and when the DVA is attaining an extremum, respectively. Once the FGS locks onto the target stars it simply attempts to minimize the root mean square (RMS) drift from these original starting positions, and thus cancels out the bulk effects of velocity aberration and DVA to first order. The ADCS cannot change the fact that the starfield appears to contract and expand over time, giving a maximum excursion for any star of about 0.6 pixels over the high-DVA quarters. To high precision, the dynamical boresight essentially follows the aberrated location of the center of the FOV. However, the clock angle of the spacecraft attitude depends on the reference frame used to establish the nominal pointing during commissioning. Once this angle was set, it, together with the spacecraft trajectory, completely determines the nominal pointing over all time. We developed a simple heuristic model for the tracking behavior of *Kepler’s* ADCS.

The SOC pointing model is calculated by solving for the spacecraft boresight $\alpha(t)$, $\delta(t)$, and roll angle, $\phi(t)$, that minimize the expected deviations of four reference points chosen on modules 7, 9, 17, and 19 that are close to the FGS sensors as functions of time, t . Empirically, there was no significant difference between minimizing the discrepancy in centroid location from the previous time step versus the first time step. This is convenient as it allows us to generate

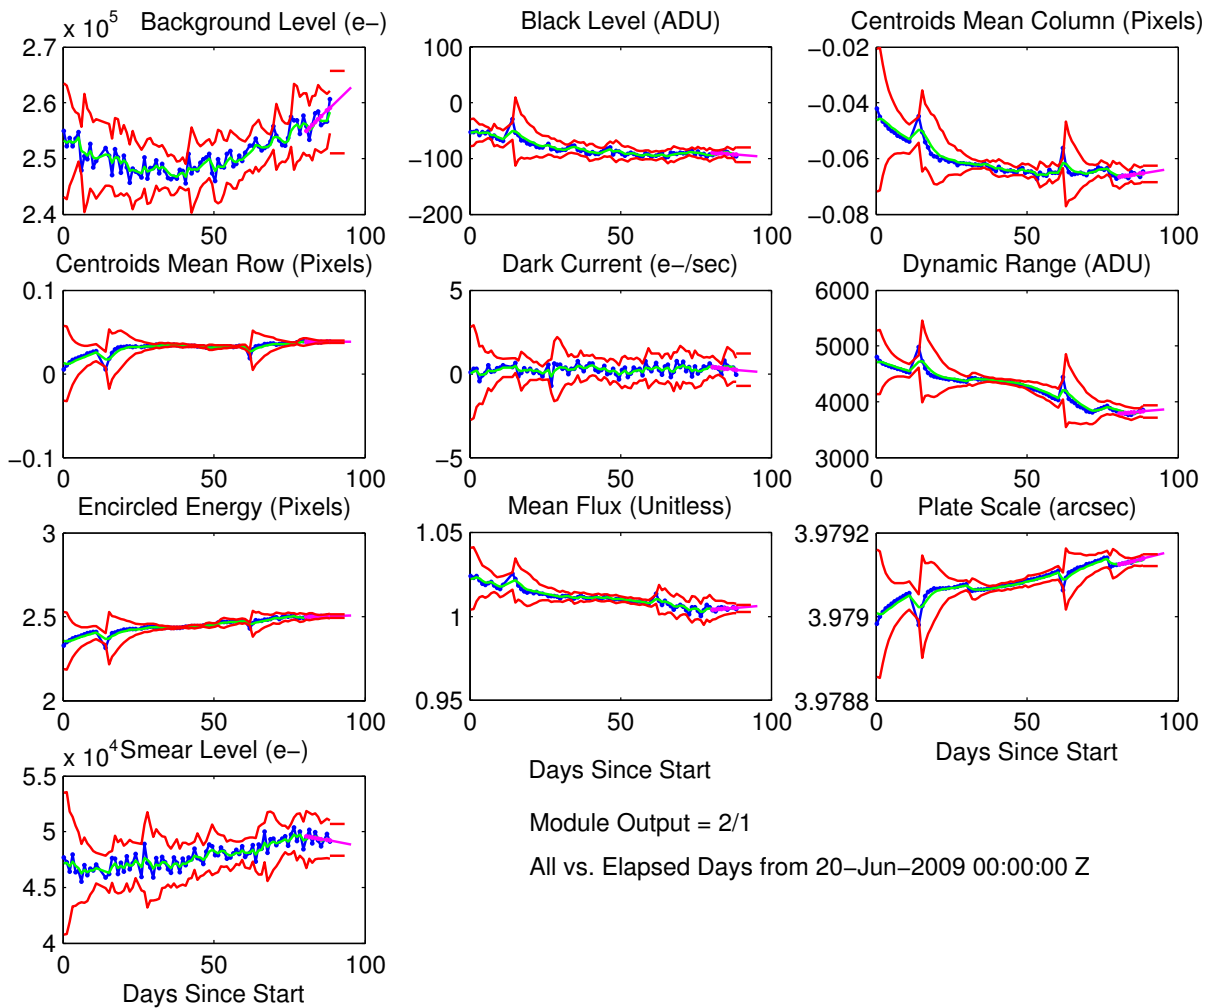


Figure 5. Tracking and trending plots for the metrics in Table 1 calculated by PDQ for each CCD channel during Q2. The blue curve is the time series for each metric. The green curve is a trend obtained by applying a moving average with an exponential forgetting factor to the metric time series. The red curves on either side of the data and smoothed trend are adaptive bounds based on estimates of the natural variation in the metric via a moving variance from the smoothed trend with an exponential forgetting factor. The magenta line is fitted to the last week of data available and projected into the future. The time required to break the current dynamic bounds in the future is called the horizon time. If the horizon time is less than one week, an alert is triggered and reported by PDQ. Any metric which breaks the dynamic or fixed bounds is reported as a yellow or red alert, respectively. Perturbations corresponding to operational events and activities, such as a safe mode event on day 15, and a monthly downlink on day 62, can be seen in several of these metrics.

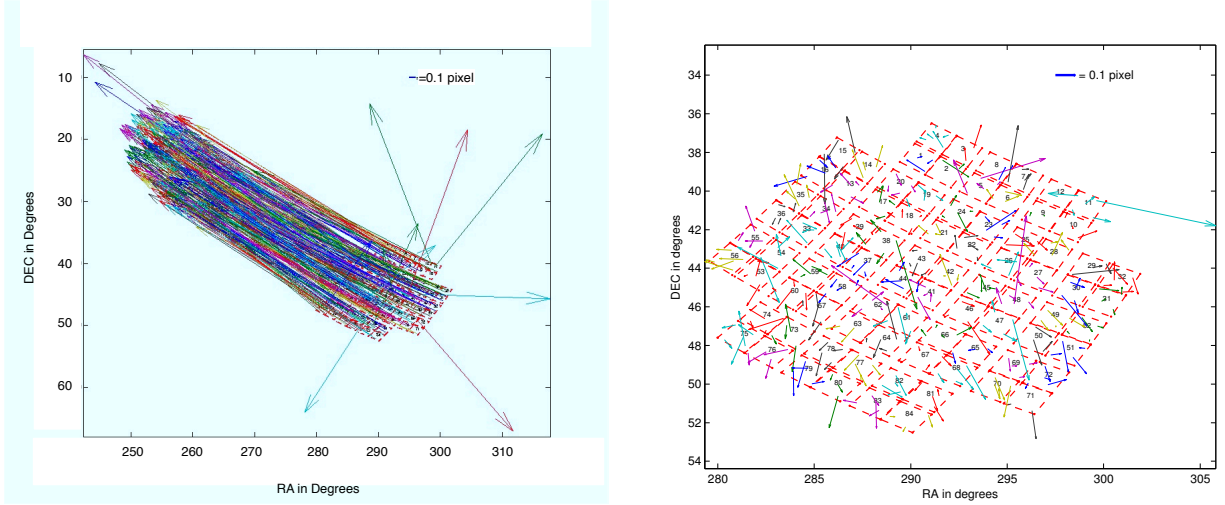


Figure 6. Quiver plots illustrating the errors between the measured locations of reference pixel target stars across the FOV both before, left panel, and after, right panel, a pointing adjustment was made based on the results of PDQ's analysis. The initial pointing error was a nearly pure lateral offset of $\sim 16''$. The outlines of the CCD readouts are indicated in the figures as well, with a scale given in degrees for the projection of the CCDs onto the sky. The quivers are scaled so that 0.1 pixels or $0.4''$ is the same length as 1° on the projected sky coordinates for the x and y axes. Note that the post-tweak residuals are uncorrelated and at a level of ~ 0.1 pixels, which is the typical size of the centroiding biases we observe in PRF-fitted centroids of bright, unsaturated stars. The attitude adjustment brought the maximum attitude residual error to $0.018''$.

the nominal pointing at any given time step rather than needing to “integrate” the pointing from the reference time step. Whenever the SOC receives a new trajectory file it updates the pointing model, predicting the nominal pointing over the entire course of the mission. The ability to make occasional adjustments to the attitude of the spacecraft allows us to make the pointing model into a self-fulfilling prophecy so that the target tables well in advance of their use are valid when they are instantiated on board the spacecraft.

4.2 To Tweak or Not to Tweak: That is the Question

When is the pointing error so egregious as to risk invalidating the stellar target definitions? The Science Office will nominally recommend an attitude adjustment if the maximum attitude residual (MAR) described in Sec. 4.4 exceeds or is expected to exceed 0.1 pixels anywhere on the focal plane. A pointing discrepancy of ≤ 0.1 pixels is acceptable since it is small compared to the maximum discrepancy due to DVA of ~ 0.6 pixels. The effects of DVA, however, can be predicted with great accuracy in advance, given the simple heliocentric orbit along which *Kepler* coasts.

There have been several occasions on which pointing tweaks were necessary. For example, after the first roll maneuver in June 2009 there was a ~ 4 pixel or $16''$ error in the pointing. Although this error was significantly larger than any that had been tested, PDQ successfully formulated a valid pointing adjustment to bring *Kepler's* boresight within 0.01 pixels or $0.04''$ of the desired location. Figure 6 shows the measured centroids for each reference pixel target star both before and after the tweak maneuver at the start of Q2. Moreover, during Q2 we observed larger than expected drift in the pointing behavior and requested pointing tweaks at some of the monthly downlink events to reset the boresight of the telescope. The ADCS algorithm parameters have since been updated and the drift has significantly diminished, so much so that we have not requested pointing tweaks at the monthlies since Q2.

4.3 Calculating the Delta Quaternion

To formulate the delta quaternion representing the change in attitude to specify to the spacecraft, we need to convert the nominal and the actual pointing into Euler 3-2-1 angles.¹³ This is accomplished by the following:

$$\begin{aligned}
 rot_3 &= \alpha \\
 rot_2 &= -\delta \\
 rot_1 &= 110^\circ + 13^\circ + 90^\circ (season - 1) + \phi + 180^\circ.
 \end{aligned} \tag{2}$$

The right ascension, α , specifies the angle of the boresight along the J2000 equator from the origin at the vernal equinox. The declination, δ , specifies the angle from the equator towards the celestial pole to the boresight. The roll angle, ϕ , together with some additional parameters, specifies the tilt of the focal plane relative to equatorial coordinates. The first term of the first line of Eq. 2, 110° , is the first roll angle, and 13° , is the nominal clock angle of the FPA relative to spacecraft coordinates. The spacecraft executes roll maneuvers every ~ 93 days to reorient the solar panels towards the Sun, so that there is a periodic 90° increment throughout the mission with the change of the spacecraft's roll orientation with $season \in \{1, 2, 3, 4\}$. Finally, the 180° accounts for the fact that the focal plane is observing an image of the FOV reflected from a mirror, so that left and right, and up and down are reversed.

The Euler 3-2-1 angles for actual and desired spacecraft attitude are then converted to quaternions $Q = q_1 \hat{i} + q_2 \hat{j} + q_3 \hat{k} + q_4$ by representing each rotation angle as a separate quaternion and then multiplying the resulting quaternions in the appropriate order:

$$\begin{aligned} Q_z &= [0, 0, \sin(\text{rot}_3/2), \cos(\text{rot}_3/2)] \\ Q_y &= [0, \sin(\text{rot}_2/2), 0, \cos(\text{rot}_2/2)] \\ Q_x &= [\sin(\text{rot}_1/2), 0, 0, \cos(\text{rot}_1/2)], \end{aligned} \quad (3)$$

where the quaternion is obtained by

$$Q_{zyx} = Q_z \otimes Q_y \otimes Q_x, \quad (4)$$

and \otimes represents quaternion multiplication. Note that the convention we've adopted is that successive rotations are represented by multiplying an existing attitude quaternion *on the right* by successive attitude adjustment quaternions. That is, $Q_B = Q_A \otimes \Delta Q$ is the attitude obtained by performing an attitude maneuver ΔQ from an initial attitude, Q_A .

The delta quaternion we desire is the rotation from the reconstructed pointing, Q_{actual} , to the desired pointing, Q_{desired} :

$$\Delta Q = Q_{\text{actual}}^* \otimes Q_{\text{desired}}, \quad (5)$$

where Q^* denotes the conjugate of Q , which is the inverse for attitude quaternions. Hence, the reconstructed attitude and the nominal attitude are related to each other by the delta quaternion through

$$Q_{\text{actual}} = Q_{\text{desired}} \otimes \Delta Q^*. \quad (6)$$

Therefore the delta quaternion that the SOC specifies to the MOC for upload to the spacecraft does not represent the attitude maneuver needed to rectify the pointing error, but rather its inverse. The ADCS presents the delta quaternion as an artificial pointing error to its Kalman filter, which dutifully slews the photometer to null out this term in addition to the real-time pointing errors it is responding to.

4.4 Maximum Attitude Residuals

From a target management perspective, the most important metric calculated by PDQ is the maximum attitude residual (MAR) metric. Once the attitude has been reconstructed, it can be used to determine the maximum discrepancy between the nominal focal plane coordinates of any target star and the actual coordinates. The MAR is inspected after each quarterly roll maneuver and after each break in science operations that requires a re-acquisition of the science attitude. If the MAR is greater than 0.1 pixels, or is likely to break this bound in the next month of operations, then the Science Office requests a tweak maneuver using the delta quaternion furnished by PDQ to bring the pointing within specifications to ensure the validity of the target tables.

Figure 7 shows the tracking and trending reported in the PDQ focal plane summary report for quarter 3.

5. CONCLUSIONS

The *Kepler* spacecraft has been orbiting the Sun in an Earth-trailing orbit since its launch on March 6, 2009. The SOC Pipeline component PDQ has been a critical component of flight operations to ensure that the spacecraft attitude is valid for science operations, and to monitor gradual changes in the spacecraft and instrument performance, such as a greater than expected pointing drift early in science operations that was mitigated by updates to the spacecraft's software parameters




Figure 7. Tracking and trending changes in the spacecraft pointing during Q3 presented in the PDQ report. PDQ reconstructs the photometer boresight right ascension (a), declination (b), and roll angle (c) and compares these to the pointing model. The maximum discrepancy (d) between the location of any target star in the FOV from its nominal location to its actual location is reported as the maximum attitude residual (MAR). If the MAR exceeds 0.1 pixels, or is likely to exceed 0.1 pixels before the next monthly contact, a pointing adjustment is requested based on PDQ's attitude solution in order to prevent the target stars from wandering outside their aperture masks. Note that the stars moved by only $\sim 0.02''$ or 0.005 pixels over the course of Q3. The different curves in each panel are defined in Fig. 5.

for the FGS centroiding algorithm. PDQ has also proven its mettle by giving a clear indication of the failure of module 3 in January 2010. As the science data accumulate for 31+ days on the SSR, PDQ allows us peace of mind that we are collecting the right data for the planetary target stars and that the photometer continues to operate in a manner consistent with obtaining the ultra-high precision required for *Kepler* to be able to discover Earth-size planets transiting Sun-like stars over its 3.5+ year mission.

ACKNOWLEDGMENTS

The authors would like to thank Dave Pletcher for his leadership of the SOC, Bill Borucki and David Koch for their leadership of the *Kepler Mission*, and Sue Blumenberg for her careful reading and editing of this manuscript.

Funding for the *Kepler Mission* is provided by NASA's Science Mission Directorate.

REFERENCES

- [1] Borucki, W., *et al.*, “*Kepler* planet-detection mission: introduction and first results,” *Science* **327**, 977–980 (2010).
- [2] Koch, D. G., *et al.*, “*Kepler Mission* design, realized photometric performance, and early science,” *ApJL* **713(2)**, L79–L86 (2010).
- [3] Van Cleve, J., and Caldwell, D. A., *Kepler Instrument Handbook, KSCI 19033-001*, NASA Ames Research Center, Moffett Field, CA (2009).
- [4] Caldwell, D. A., *et al.*, “Instrument performance in *Kepler*'s first months,” *ApJL* **713(2)**, L92–L96 (2010).
- [5] Middour, C., *et al.*, “*Kepler* Science Operations Center architecture,” *Proc. SPIE* **7740**, in press (2010).
- [6] Li, J., *et al.*, “Photometer performance assessment in *Kepler* science data processing,” *Proc. SPIE* **7740**, in press (2010).
- [7] Haas, M., *et al.*, “*Kepler* science operations,” *ApJL* **713(2)**, L115–L119 (2010).
- [8] Bryson, S. T., *et al.* “Selecting pixels for *Kepler* downlink,” *Proc. SPIE* **7740**, in press (2010).
- [9] Quintana, E. M., *et al.*, “Pixel-level calibration in the *Kepler* Science Operations Center pipeline,” *Proc. SPIE* **7740**, in press (2010).
- [10] Allen, C. A., *et al.*, “*Kepler Mission*'s focal plane characterization models implementation,” *Proc. SPIE* **7740**, in press (2010).
- [11] Bryson, S. T., *et al.* “The *Kepler* pixel response function,” *ApJL* **713(2)**, L97–L102 (2010).
- [12] Batalha, N. M., *et al.*, “Selection, Prioritization, and Characteristics of *Kepler* Target Stars,” *ApJL* **713(2)**, L109–L114 (2010).
- [13] Kuipers, J. B., *Quaternions and Rotation Sequences: A Primer with Applications to Orbits, Aerospace, and Virtual Reality*, Princeton University Press (1999).