Processing CCD images to detect transits of Earth-sized planets: Maximizing sensitivity while achieving reasonable downlink requirements

Jon M. Jenkins[a], Fred Witteborn[b], David G. Koch[c], Edward Dunham[d], William J. Borucki[c], Todd F. Updike[e], Mark A. Skinner[e], and Steve P. Jordan[e]

[a] SETI Institute, MS 245-3, Moffett Field, CA 94035
[b] Orbital Sciences Corp, MS 245-6, Moffett Field, CA 94035
[c] NASA Ames Research Center, MS 245-6, Moffett Field, CA 94035
[d] Lowell Observatory, 1400 W. Mars Hill Rd, Flagstaff, AZ 86001
[e] Ball Technologies Corporation, 1600 Commerce St., Boulder, CO 80301

ABSTRACT

We have performed end-to-end laboratory and numerical simulations to demonstrate the capability of differential photometry under realistic operating conditions to detect transits of Earth-sized planets orbiting solar-like stars. Data acquisition and processing were conducted using the same methods planned for the proposed *Kepler Mission*. These included performing aperture photometry on large-format CCD images of an artificial star field obtained without a shutter at a readout rate of 1 megapixel/sec, detecting and removing cosmic rays from individual exposures and making the necessary corrections for nonlinearity and shutterless operation in the absence of darks. We will discuss the image processing tasks performed "on-board" the simulated spacecraft, which yielded raw photometry and ancillary data used to monitor and correct for systematic effects, and the data processing and analysis tasks conducted to obtain lightcurves from the raw data and characterize the detectability of transits. The laboratory results are discussed along with the results of a numerical simulation carried out in parallel with the laboratory simulation. These two simulations demonstrate that a system-level differential photometric precision of 10^{-5} on five-hour intervals can be achieved under realistic conditions.

Keywords: extrasolar planets, Earth-size planets, photometry, CCDs, space mission, planetary transits

1. INTRODUCTION

The proposed *Kepler Mission* seeks to determine the frequency of Earth-sized planets by transit photometry. *Kepler* will stare at the same 84 deg^2 field of view (FOV) for four years, conducting differential aperture photometry to detect planets with orbital periods up to two years[1]. The desire to study ~170,000 stars simultaneously by CCD photometry to detect transits of planets as small as Earth imposes significant restrictions on the amount of data that can be downlinked from the spacecraft. In particular, the downlink rate will not permit continuous transmission of entire CCD frames. However, the data link will permit transmission of pixels in the photometric apertures of the target stars without the use of lossy compression algorithms.

We have conducted a series of end-to-end laboratory and numerical simulations to demonstrate the ability to repeatedly and reliably detect fractional changes in stellar brightness of 8×10^{-5} for stars brighter than $m_v=12$. Both the hardware and software simulations include realistic noise sources such as bright 4th magnitude stars, shutterless readout, fast optics, cosmic rays and spacecraft pointing jitter. The companion paper by Koch et al. describes the laboratory testbed in detail and presents the results of various tests conducted therein[2]. In this paper, we describe the algorithms used to achieve the desired precision, including novel techniques for performing aperture photometry to obtain raw stellar fluxes and for removing systematic effects from the resulting lightcurves.

The paper is organized as follows. First, we give a description of the image processing tasks that are to be conducted onboard the spacecraft. Next, we detail the method by which the pixel weights for conducting the aperture photometry are designed, along with a discussion of pre-masking pixels for flux calculations. Third, we present the processing of the raw fluxes to produce relative fluxes and the method of removal of systematic errors. Fourth, we discuss the numerical simulations carried out in parallel with the laboratory work. Fifth, we examine the data volume likely to be generated by the *Kepler Mission*.

Finally, we provide conclusions about the applicability of the techniques described herein to current and future ground-based and space-flown photometry missions.

2. ONBOARD PROCESSING OF CCD FRAMES

The *Kepler* Testbed represents an end-to-end test of the *Kepler Mission* concept. As such, the hardware and software were designed to test key aspects of the mission. All software tasks described in this section can be implemented on the current *Kepler* spacecraft design. However, some of the tasks may not be required with the actual flight hardware. For example, the nonlinearity of the CCD output electronics (§2.2) is much higher than that expected for the spacecraft electronics. Hence, a nonlinearity correction may not be required onboard the spacecraft, decreasing hardware requirements. The corrections for bias and shutterless readout may be deferred to ground processing, if the increase in downlink requirements is warranted by the decrease in onboard processing capabilities. These two tasks are not numerically intensive, however, so that we expect to perform them onboard. Naturally, the software evolved during the course of conducting the laboratory demonstration. This discussion is confined to the final algorithms that were used to analyze the laboratory data. The following describes the tasks that nominally will be performed aboard the *Kepler* spacecraft during its mission as they were applied in the laboratory demonstration. These steps include rejection of cosmic ray events, co-adding of 2.5-s exposures to obtain 3- or 15-min integrated frames, and correction of star pixels for CCD bias, nonlinearity and shutterless readout.

2.1. Cosmic ray removal and co-adding CCD frames
The magnitude range of our target stars is from m_v=9 to m_v=14. Saturation for m_v<9 for the *Kepler* photometer sets the exposure time for individual frames at 2.5 s, with readouts of 0.5 s. Each individual exposure is examined on readout for cosmic ray events and then added to a running sum array. Pixels flagged for cosmic ray events are not added to the running sum, and the cosmic ray counter for each of these pixels is incremented by one for each cosmic ray event detected. In the laboratory demonstration, the limitations of the available hardware (a Sparc5 workstation running at 170 MHz) prevented us from applying the proposed cosmic ray detection algorithm in real time. The actual algorithm will examine the change in pixel brightness from a simple predictive estimate for each pixel based on previous valid samples. Instead, we used a simple fixed threshold to detect cosmic rays in real time. This algorithm worked well for cases without spacecraft pointing jitter, but was not effective when spacecraft jitter was included. On the other hand, cases for which simulated cosmic rays were injected but not removed were not significantly affected by the cosmic rays. The anticipated cosmic ray flux rate implies that each pixel will be hit about once every day[2]. Based on the laboratory data, cosmic ray rejection may not be required aboard the spacecraft, unless the cosmic ray flux or average energy deposited are significantly different from the expected values. Aside from co-adding individual frames in memory, cosmic ray rejection is the only task that must be conducted on 2.5-s frames. All other operations are conducted on 3-min or 15-min frames. The co-adding process is continued either for 60 or 300 times, resulting in a frame representing 3 or 15 minutes of exposure at a duty cycle of 83.3%. After completion of co-addition, each pixel is corrected for the number of cosmic rays detected by scaling the summed pixel value by the total number of individual exposures divided by the number of exposures not flagged for cosmic ray events. This effectively replaces the missing pixel measurements with the average flux received by each pixel over the 3- or 15-minute period during the unflagged exposures. (For example, a pixel flagged for one cosmic ray event over 300 individual exposures would be scaled by 300/299=1.003344.) In the laboratory demonstration, the corrected frames are written to disk as FITS files for later processing.

2.2. Bias and non-linearity corrections

Figure 1 presents a typical CCD frame for the *Kepler* testbed CCD with the bias strip and readout (smear) correction region indicated by superimposed boxes. The readout amplifiers are located at the upper corners of this figure, so that pixels are read out horizontally from the outermost pixels inward to the center of the CCD after being clocked to the top row. In the laboratory demonstration the bias is estimated for each row by overscanning both amplifiers on the CCD by 12 pixels after all pixels on the current row have been read out. On each summed frame, the 12 bias values for each row are averaged, and then a 51-point boxcar filter is applied to the resulting column vector to further reduce the noise in the bias measurement. The filtered bias estimate for each row is subtracted from corresponding pixels of interest, including star pixels, and pixels in the readout smear correction region.

Following the bias correction, the star pixels are corrected for nonlinearity. As shown in figure 2, the nonlinearity correction rises linearly from 1.0 at 0 ADU to 1.07 at 3.5×10^4 ADU (about 50% of full well), flattens out and then rises again to ~1.10 at

full well. Note that 1 ADU = 11.1 e⁻. This nonlinear response is due almost entirely to the fast readout rate employed in the demonstration. No significant nonlinearity above 1% was seen at lower readout rates.

2.3. Shutterless readout correction

After the nonlinearity correction, the star pixels are corrected for the effect of the shutterless readout. Here we note that during readout, each pixel "sees" 0.5/1024=0.5 msec of flux from each pixel above it as the CCD is clocked out. Prior to the 2.5-s integration, however, this pixel was clocked into position on the previous readout, and "saw" 0.5 msec of flux from each row below its nominal position. Thus, except for 0.02% of their own fluxes during the 2.5-s exposure, all pixels in a column pick up the same excess flux from other pixels in its column during readout. To estimate this readout smear flux, a 50-row deep, non-illuminated region below the stellar images is examined. The 50 bias-corrected values in each column of this region are averaged and then subtracted from the star pixels in the corresponding columns.
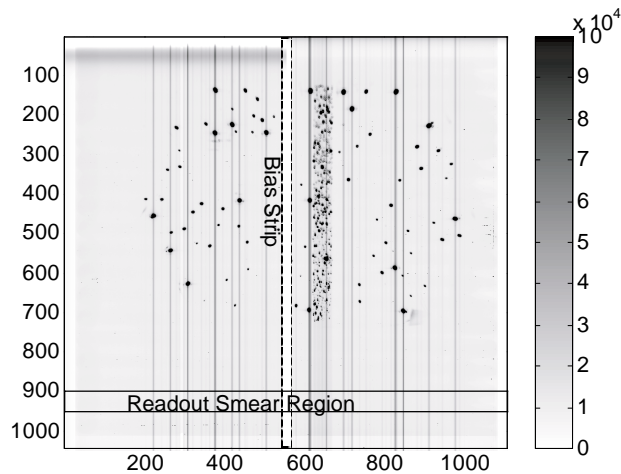

Figure 1. A typical 3-min CCD frame with bias removed clipped at 10⁵ ADU illustrating the bias strip and the readout smear calculation strip. The grayscale units are in ADU.
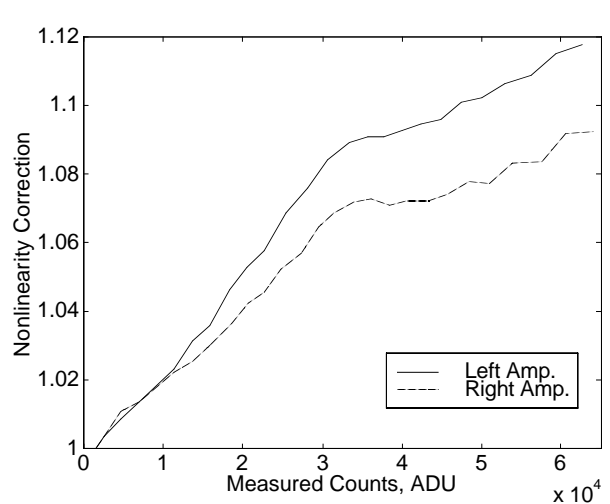
Figure 2. Nonlinearity response of the CCD and its readout electronics for each amplifier.

2.4. Data extraction and aperture photometry

Following the above corrections at the end of each 3- or 15-min integration, an 11 by 11 pixel region containing each target star is written to disk in the lab for later processing. In addition, centroids of all target stars are computed and saved to disk. Aperture photometry is applied to star pixels to obtain raw flux time series. The aperture photometry performed on these laboratory data is different from that usually employed. It consists of estimating the flux of each target star as the weighted sum of pixels in its photometric aperture at a given time step, rather than by simply summing the fluxes from some subset of pixels in the 11 by 11 aperture. The weights are chosen in such a way as to minimize the sensitivity of the flux estimate to image motion (compounded by pixel-to-pixel variations) and to changes in the point spread function (PSF). This is similar to PSF-fitting photometry, in which the flux at each time step is estimated by a weighted sum of the pixels in the fitting aperture, where the weights depend on the PSF, the background noise and the image position at that particular time. In our case, however, the weights have fixed values, and function to reduce the sensitivity to image changes, not to correct for them. The next section derives an expression for the optimal pixel weighting performed in the laboratory experiment and discusses implementation issues.

3. OPTIMAL PIXEL WEIGHTING

Motions of stellar images over a finite photometric aperture cause apparent brightness changes (even with no intra- or inter-pixel sensitivity variations). The wings of any realistic PSF cause these motion-induced brightness variations, as they extend outside of any reasonable photometric aperture. Pixel-to-pixel variations generally exacerbate motion-induced brightness variations as well as causing apparent changes in the PSF. In addition, changes in platescale and focus also induce apparent brightness changes in measured stellar fluxes. Figure 2 of Koch et al. presents an example from the *Kepler* Testbed[2]. Several

possible remedies to these problems exist: 1) Calibrate the response of each star's measured brightness as a function of position and focus and use this information to correct the measured pixel fluxes. 2) Regress the lightcurves against the measured motion and focus or other correlated quantity to remove these effects. 3) Calculate the stellar fluxes using weighted sums of the aperture pixels in such a way as to reduce the sensitivity to the typical image motion and focus changes.

The first solution requires detailed knowledge of the 3-D correction function for each star, and must be applied on timescales short enough so that the change in position and focus is small compared to the full range of motion and focus change. This solution is equivalent to PSF-fitting photometry. For the *Kepler Mission*, the attitude control system operates on timescales much shorter than 15 minutes, so that the motion becomes decorrelated after about 25 seconds. One component of focus and platescale change will be long term, due to the apparent 1° per day rotation of the sun about the spacecraft[1]. Changes on timescales this long can be neglected for the purposes of transit photometry, so long as the amplitudes are not large enough to move the target stars by significant fractions of a pixel. The short coherence time of the spacecraft jitter would necessitate the application of the flux corrections after one or several readouts, which is impractical. The second solution has been previously demonstrated in obtaining 1x10[-5] photometry for front-illuminated CCDs[3] and for back-illuminated CCDs[4]. Our modification of this method is presented in §4. In contrast to the first approach, the third solution is feasible if the image motion and focus changes are approximately wide-sense stationary random processes (i.e. the statistical distributions of changes in position and focus are constant in time[5]). Strict wide-sense stationarity is not required, however, it simplifies the implementation of the method, as updates in the pixel weights would not be required in between spacecraft rotations (which occur every 3 months). What remains is the problem of designing the pixel weights themselves. The remainder of this section follows the derivation of a formula for obtaining the optimal pixel weights, and gives examples of their effectiveness in reducing sensitivity to motion.

3.1. Theoretical development

We wish to derive an expression for the optimal pixel weights minimizing the combination of sensitivity of the flux of a star to image change and the effects of the weights on shot noise and background noises. This approach is motivated by a signal processing perspective in which aperture photometry is viewed as applying a finite-impulse response (FIR) filter to a temporally-varying 2-D waveform. In the 1-D signal-processing analog, the desire is to shape the frequency response of a FIR filter to reduce or enhance the sensitivity of the filter to energy in certain wavelength regions. In the problem at hand, the desire is to use the free parameters available (the pixel weights) to minimize the response in the flux measurement to image motion and PSF changes. The following assumptions are made: 1) the PSF and its response to image motion and focus change are well-characterized, 2) the distribution of background stars is known, and 3) the background noise sources are well-characterized. Consider a set of N images of a single target star and nearby background stars consisting of M pixels ordered typographically (i.e. numbered arbitrarily from 1 to M). Assume that the range of motion and focus change over the data set are representative of the full range of motion and focus changes. Define the error function, E, as the combination of the mean fractional variance between the pixel-weighted flux and the mean, unweighted flux and a second term accounting for the effect of the pixel weights on the shot and background noise:

$$E \equiv \frac{1}{N}\frac{1}{\overline{B}^2}\sum_{n=1,\ldots,N}\left(\overline{B}-\sum_{j=1,\ldots,M}w_j b_{n,j}\right)^2 + \frac{\lambda}{\overline{B}^2}\sum_{j=1,\ldots,M}w_j^2\left(\overline{b}_j+\sigma_j^2\right), \qquad (1)$$

where

$$b_{n,j} \equiv \text{jth pixel value at timestep } n, \, n=1,\ldots,N, \, j=1,\ldots,M$$

$$w_j \equiv \text{weight for pixel j}$$

$$\overline{b}_j \equiv \text{mean pixel value for pixel j} \qquad (2)$$

$$\overline{B} \equiv \text{mean flux with all weights set to 1}$$

$$\sigma_j^2 \equiv \text{background noise variance for pixel j,}$$

and all quantities are expressed in e⁻. Here we take the shot noise to be due entirely to the star itself, and the background noise to be a zero-mean process which includes such noise sources as readout noise and dark current. This implies that the images

have been corrected for all non-zero-mean noise sources such as dark current and the readout smear flux. We further assume that the background noise sources are uncorrelated from pixel to pixel. If this is not the case, the second term of (1) can be augmented to account for the correlation. The scalar $\lambda \in [0, \infty)$ determines the balance between the desire to minimize the difference between the flux estimate and the mean flux value, and the desire to minimize the accompanying shot noise and the background noise. For this situation, we would normally set $\lambda = 1$.

The error function in (1) is quadratic, and therefore admits a closed-form solution in matrix form:

$$\mathbf{w} = \left[\tfrac{1}{N}\mathbf{B}^T \cdot \mathbf{B} + \lambda \mathbf{D}\right]^{-1} \cdot \overline{\mathbf{b}}\,\overline{B}, \qquad (3)$$

where

$$\mathbf{B} \equiv \left\{b_{n,j}\right\}, \, n = 1,\ldots,N; \, j = 1,\ldots,M$$

$$\mathbf{D} \equiv \left\{D_{i,j}\right\} = \begin{cases} \overline{b}_i + \sigma_i^2, \, i = j = 1,\ldots,M \\ 0, \qquad i \neq j \end{cases} \qquad (4)$$

$$and$$

$$\overline{\mathbf{b}} \equiv \left\{\overline{b}_j\right\}, \, j = 1,\ldots,M.$$

Throughout this paper, boldface symbols represent column vector or matrix quantities. For real data with noise-corrupted images, the scalar λ should be adjusted to prevent over-fitting. If enough data is available, λ will be essentially 0. An alternative iterative scheme can be used that is based on the popular NLMS (normalized least mean square error) algorithm for adaptive filtering[5]. The chief advantage of such an algorithm is that the pixel weights can be designed "in place", and can be updated as necessary. This algorithm adjusts the pixel weight vector by an incremental vector opposite the direction of an estimate of the gradient of the error function.

Taking the expression

$$\hat{E}(n) = \left(\overline{B} - \mathbf{b}_n^T \cdot \mathbf{w}_n\right)^2 = \left[\overline{B} - \left(\overline{\mathbf{b}} + \Delta\mathbf{b}_n^T\right)\cdot \mathbf{w}_n\right]^2 = \left(\overline{B} - \overline{\mathbf{b}}\cdot\mathbf{w}_n - \Delta\mathbf{b}_n^T\cdot\mathbf{w}_n\right)^2 = \left(\Delta\mathbf{b}_n^T\cdot\mathbf{w}_n\right)^2 \qquad (5)$$

as the error estimate at time n, where $\Delta\mathbf{b}_n$ is the difference between the average pixel fluxes and those at the n^{th} time step, the update to the weight vector at time step n is given by

$$\mathbf{w}_n = \mathbf{w}_{n-1} - \mu \frac{\nabla E(n)}{\Delta\mathbf{b}_n^T \cdot \Delta\mathbf{b}_n + \varepsilon} = \mathbf{w}_{n-1} - \mu \frac{\Delta\mathbf{b}_n^T \cdot \mathbf{w}_{n-1}}{\Delta\mathbf{b}_n^T \cdot \Delta\mathbf{b}_n + \varepsilon}\Delta\mathbf{b}_n^T, \qquad (6)$$

where μ is a positive scalar that controls the rate of adaptation (and convergence in the case of stationary noise and motion) and ε is a small positive number to ensure stability in cases where the instantaneous error is 0. Note that the term for shot and background noise does not appear here. It is not necessary as μ can be adjusted to prevent over-fitting the noise, and the algorithm is mainly of interest in the case of noise-corrupted images.

In terms of implementation on the *Kepler* spacecraft, equation (3) is preferred, as it reduces the computations required onboard. This approach would require the collection of adequate samples of star pixels to recover well-sampled (super-resolution) scenes for each target star. This might be avoided with proper calibration of the optics and CCDs along with a high-resolution catalog of stars in the FOV. If necessary, the adaptive scheme of (6) could be implemented, with proper choice of scheduling and for the value of μ to insure that the adaptation takes place over time scales significantly longer than a transit.
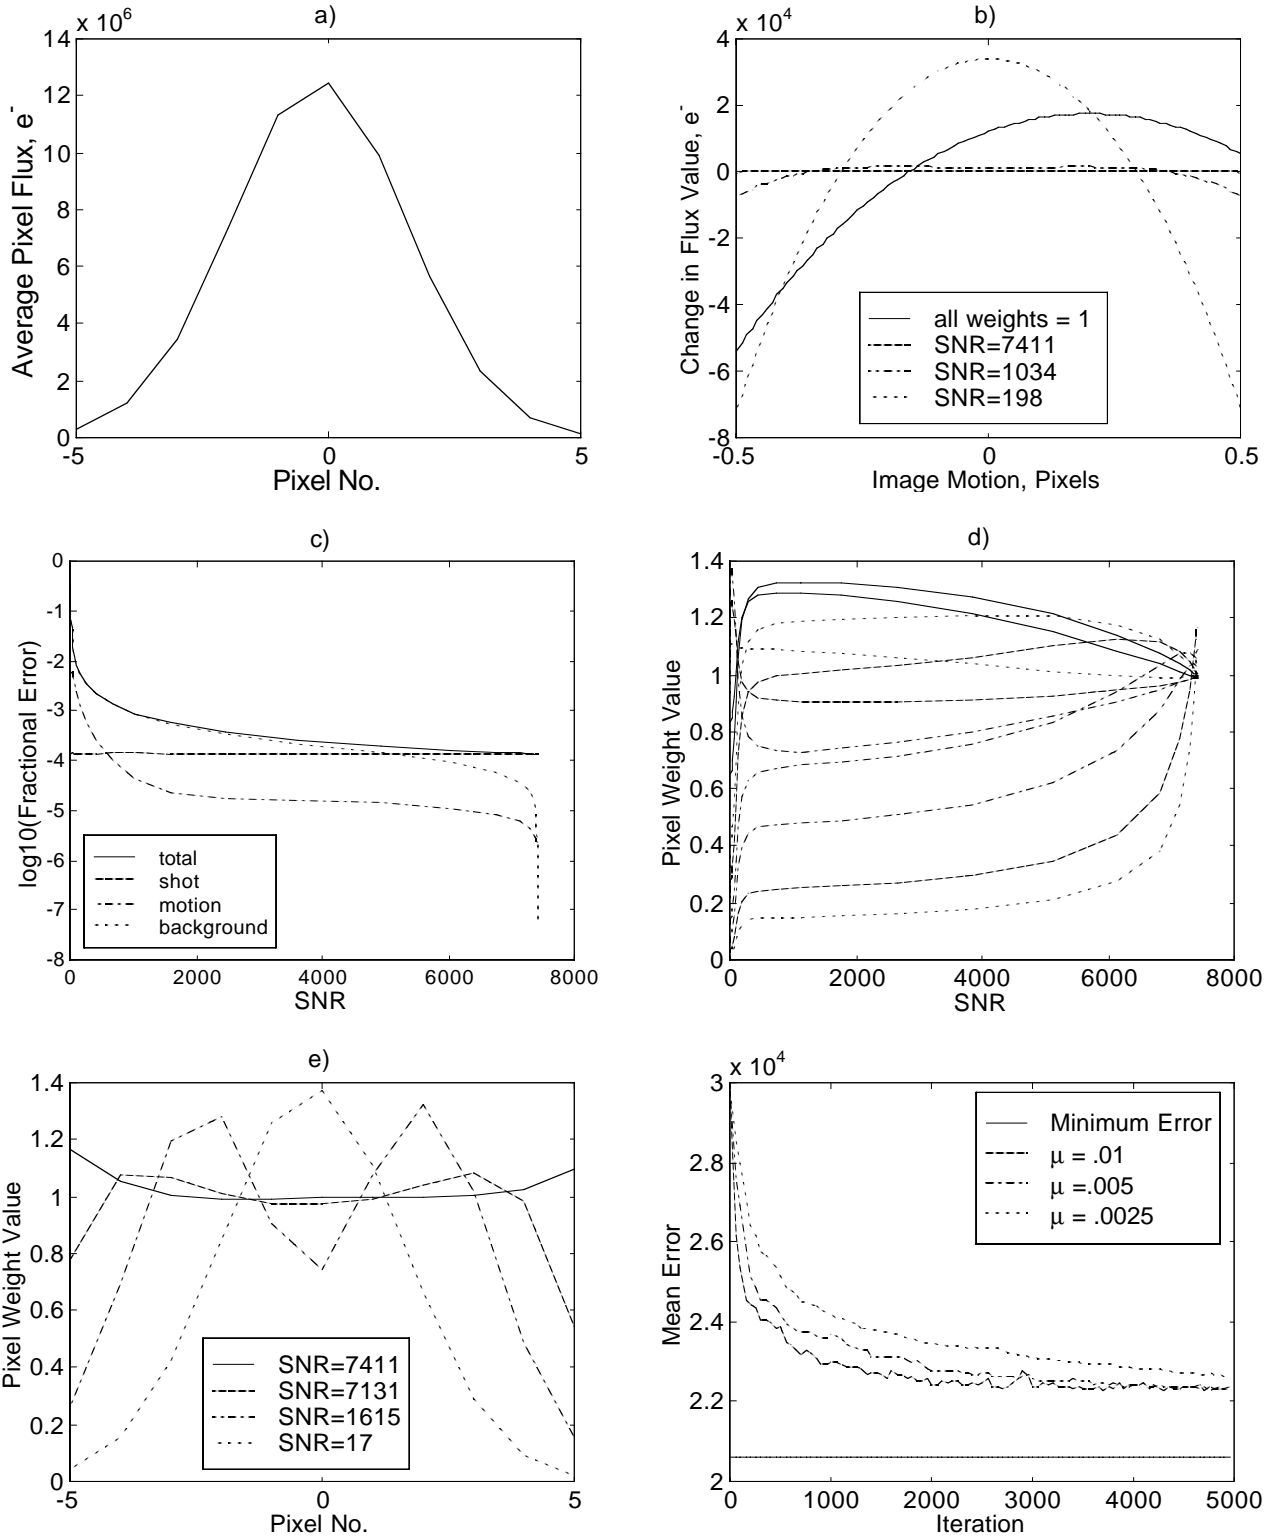
Figure 3. Results from a numerical simulation illustrating various properties of optimal pixel weights as a function of SNR. See text for a detailed explanation.

3.2. A one-dimensional example

In this section we provide a 1-D example to examine various properties of optimal pixel weighting. Fig 3a shows the average "image" of a Gaussian with a full width half max (FWHM) of 4 pixels on an 11 pixel aperture, nominally centered at –0.2 pixels. The integration time corresponded to 3 minutes for a 12th magnitude star for the *Kepler* photometer, yielding an average flux of 5.5×10^7 e$^-$ at each timestep. This PSF was moved over a 101-point grid in space of ± 0.5 pixels from its nominal location at -0.2 pixels and integrated over each pixel in the aperture to form a set of images. Figure 3b illustrates the response of the flux signal to image motion over the data set for optimal pixel weights corresponding to various signal-to-noise ratios (SNRs) and for unweighted aperture photometry. Here, SNR is defined as the ratio of the mean flux to the root sum square (RSS) combination of shot noise and background noise. Note that the full range of brightness variations is 1.3% for unweighted pixels, and that this is reduced to 1×10^{-5} at an SNR of 7,411. At low SNRs, the background noise dominates, and the pixel weights adjust to minimize the increased noise due to background, rather than to motion. However, the response to motion is made symmetric by the pixel weights even at low SNR so that the motion will more easily average out over timescales much longer than the coherence scale of the motion. Panel 3c presents the total expected fractional error and its three components, shot noise, motion error, and background noise, as functions of SNR. The pixel weights confine the motion error to well below the unweighted case over the range of SNRs presented here. Panel 3d shows the evolution of the optimal pixel weights as a function of SNR, while panel 3e shows the profiles of the pixel weights at four different SNR values. As the SNR deteriorates, the profile of the optimal pixel weights looks more and more like the original star profile. The final panel 3f illustrates the application of (6) to an online adaptive solution for the pixel weights. A total of 5,000 images along with shot noise and background noise of 6310 e$^-$/pix were presented to the algorithm, which was initialized with all weights equal to 1. The pixel weights converge after a few thousand iterations, corresponding to a few days of adaptation. Better initialization would result in faster convergence. We note that the excess error, or misadjustment of the weights is rather small, about 10% of the theoretical minimum error. Once convergence is achieved, μ can be reduced so that the algorithm tracks changes in the mean image position and PSF shape over timescales much longer than transits, so that transits are preserved in the resulting flux time series

3.3. Selecting pixel masks

As in conventional differential aperture photometry, optimal pixel weighting benefits from pre-masking of the pixels containing each star image to consider only those pixels with significant stellar flux content. The advantages are two-fold. First, design of optimal pixel weights for dim pixels from actual images is problematic whether the weights are generated using Eq. (3), or whether an adaptive algorithm is applied online. A great deal of data is required to reduce the uncertainties in the corresponding pixel weights to acceptable levels. Second, reducing the number of pixels for which weights are sought reduces the amount of data required for the pixel weight design. Various schemes for identifying the photometric aperture diameter and shape appear in the literature[6,7]. Here we present the method applied to the laboratory data to identify pixels allowed to participate in flux measurement for each star. The pixels are first listed in order of descending brightness. Next the cumulative SNR for the pixel list is calculated according to:
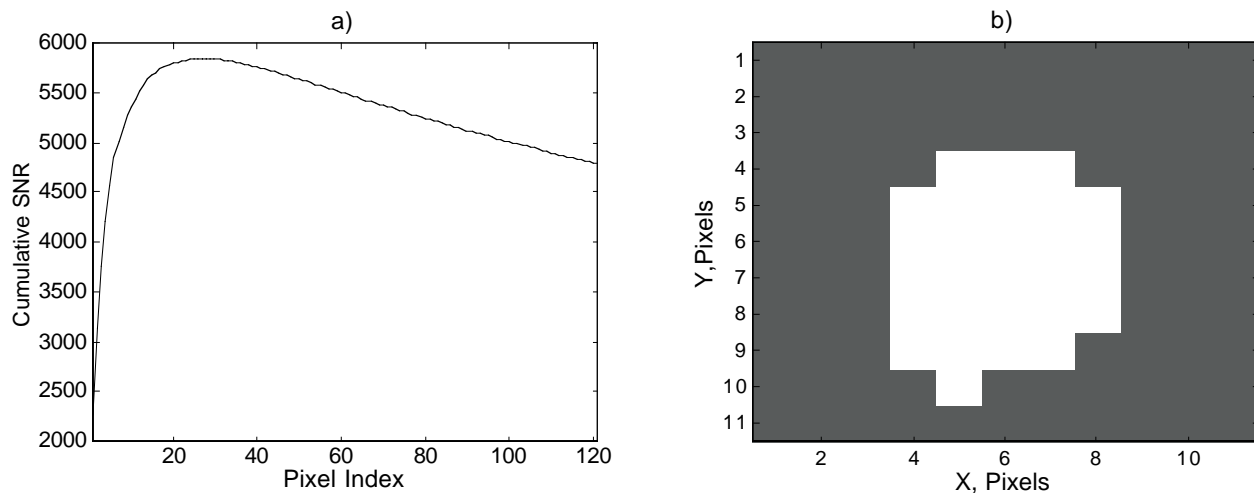

Figure 4. Cumulative SNR for star S12d (panel a) shows that only 32 pixels contribute meaningful information about the star's flux. The pixels used for flux calculations for this star are shown in white in panel b.

$$SNR(i) = \sum_{j=1,\ldots,i} \bar{b}_i \left/ \sqrt{\sum_{j=1,\ldots,i} \bar{b}_i + i\,\sigma_j^2}\right. ,\ i = 1,\ldots,M, \qquad (7)$$

where all units are in e⁻. The function $SNR(i)$ will increase as more pixels are added until the point at which the pixels are so dim as to detract from the information content of the flux estimate. All pixels beyond the point at which the maximum is attained are masked out. Figure 4a shows the cumulative SNR for star S12d, a 12[th] magnitude smear star in the laboratory demonstration. $SNR(i)$ peaks at 32 pixels. Panel 4b shows the pixels selected for inclusion in the flux estimate for this star in white, while masked-out pixels are in black. The border of the mask is roughly circular, as expected.

4. DECORRELATING NORMALIZED RELATIVE FLUXES

The optimal pixel weighting technique of §3 is effective in reducing the sensitivity of stellar flux measurements to image changes, but it cannot correct for the effects of slowly-varying systematic errors such as drift in the image position over timescales comparable to a transit. Robinson et al. (1995) solved this problem by regression of the normalized relative lightcurves in terms of image position time series[3]. Jenkins et al. (1997) included changes in PSF width along with motion in the set of regression variables[4]. Essentially, this approach fits each lightcurve in terms of a set of dependent basis functions and removes any features that are correlated with the regression data. From a linear algebra standpoint, the lightcurves and the regression data are viewed as elements of a vector space. The projection of each lightcurve onto the subspace spanned by the regression vectors is removed, yielding residual lightcurves that cannot be represented as linear combinations of the regression subspace.

In the *Kepler* testbed, the motions of the starfield images over the course of an experiment were rather complex. The dominant component was translational, caused by thermal expansion and contraction of materials in the mount, optics and structure. Differential motion of the stars also occurred, and was correlated with changes in focus and plate scale with temperature changes. Regressing on motion did not significantly improve the results. Some of these effects caused small changes in PSF shape and size as well as image position. The lightcurves obtained to this point in processing appeared to be highly correlated with respect to the long-term temporal structures so evident. Figure 2 of Koch et al. examines some of these effects[2]. Assuming that these temporal variations in apparent relative flux were due to systematic errors common to all stars, we attempted to remove them by regressing each star's lightcurve against the set of all other lightcurves. This would effectively decorrelate the resulting set of lightcurves, preserving only those signals unique to each star, such as transits. The results were highly satisfactory. The algorithm applied to obtain the K residual lightcurves is as follows, and parallels the mathematical development for optimal pixel weights:

$$\tilde{\mathbf{f}}_k = \mathbf{f}_k - \mathbf{F}_k \cdot \mathbf{c}_k \qquad (6)$$

where $\tilde{\mathbf{f}}_k$ is the k[th] residual lightcurve, $\mathbf{f}_k$ is the k[th] original normalized relative lightcurve (i.e., the relative flux time series has been divided by its mean and then added to –1 to obtain a zero-mean time series), $\mathbf{F}_k$ is a matrix whose columns contain all lightcurves other than $\mathbf{f}_k$, and $\mathbf{c}_k$ is the coefficient vector for the k[th] lightcurve. The solution for $\mathbf{c}_k$ is obtained by

$$\mathbf{c}_k = \left[\mathbf{F}_k^T \cdot \mathbf{F}_k + \lambda \mathbf{A}_k\right]^{-1} \cdot \left(\mathbf{F}_k^T \cdot \mathbf{f}_k\right), \qquad (7)$$

where

$$\mathbf{F}_k \equiv \left\{f_{n,j}\right\},\ n = 1,\ldots,N;\ j \in \{1,\ldots,K\} \setminus \{k\}$$

$$\mathbf{A}_k \equiv \left\{A_{i,j}\right\} = \begin{cases} \bar{B}_i^{-1/2}, & i = j \in \{1,\ldots,K\} \setminus \{k\} \\ 0, & i \neq j \end{cases} \qquad (8)$$

and

$$\mathbf{f}_k \equiv \left\{f_{n,k}\right\},\ n = 1,\ldots,N.$$

Here, $\overline{B}_i^{-1/2}$ is the shot noise for the ith lightcurve and λ is, as before, an adjustable non-negative scalar that is adjusted to prevent over-fitting. Some cautions are in order. We note that the number of time steps must be much larger than the number of stars used in the regression (i.e. $N >> K$). If this is not the case, random fluctuations in the lightcurves will guarantee near-full rank for the matrix $\mathbf{F}_k^T \cdot \mathbf{F}_k$, and the resulting residuals will be meaningless. If a star is a member of the set of reference stars, it cannot be regressed against a set of lightcurves containing the other normalizing stars. This is because the normalization forces the set of reference stars' lightcurves to be linearly dependent. In these cases, we have removed all reference stars from consideration in decorrelating other reference stars. One alternative solution would be to normalize each reference star by the sum of reference fluxes less their own flux. Another solution might be to perform the regression in the original space of raw flux time series on a logarithmic scale. In numerical experiments we found that this latter procedure yields results very similar to the results of the algorithm described above, except that it appeared to be more sensitive to outliers. Given this observation, decorrelating relative lightcurves as presented here is equivalent to normalizing each star's flux time series by a product of the other flux time series each raised to some power (given by the resulting regression coefficients).
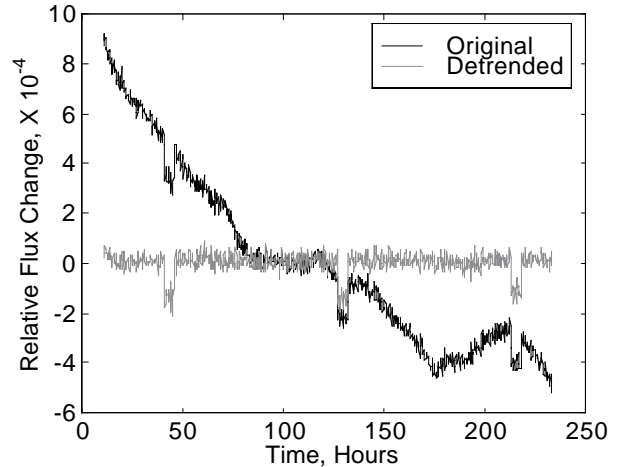

Figure 5. The result of decorrelating a star's lightcurve is shown, exhibiting the desirable feature of preservation of the compact transit features.

We conclude this section with figure 5, which illustrates the effectiveness of decorrelation in reducing undesirable long-term systematic variations in apparent brightness while preserving features such as transits. The lightcurve is for a 9th magnitude star in the *Kepler* testbed undergoing 3 transits in a 240-hour period. The resulting residual lightcurve is quite flat outside of the three five-hour, 1-Earth-Area transits. Decorrelation has effectively "whitened" the observation noise, simplifying the task of detecting the transits.

5. THE NUMERICAL SIMULATION

A numerical simulation of the flight photometer was carried out by Ball Aerospace Corporation in parallel with the laboratory simulation conducted at NASA Ames Research Center. Numerical simulation of the flight photometer allows a wider range of parameters to be explored quickly. The actual PSF of the flight 1-m Schmidt optical system can be used, or variations of the PSF can be explored. Any brightness of star and any depth or duration of transit can be simulated. Error sources can be increased in magnitude arbitrarily to determine sensitivity to errors and combinations of error sources. End-of-life design values can be compared to beginning-of-life values. Different photometric apertures can be used, and different weighting schemes can be evaluated for effectiveness. It is also possible to evaluate different CCD designs without purchasing and installing actual devices in the laboratory simulation set up. While the numerical simulation work was primarily focused on the flight hardware, we did perform a simulation of the laboratory hardware to verify the validity of the numerical simulation code.

The numerical simulation was originally built in 1994, in Ball's Integrated Modeling and Simulation Lab, and was significantly enhanced for this work. The simulation consists of a suite of numerical models that are applied manually. First, a thermal model of the photometer was built, including structure, optics, focal plane, thermal control surfaces, power, thermal links, and boundary conditions for the heliospheric drift-away orbit proposed for *Kepler*. The model is implemented in the Themal Analysis Kit (TAK) III with radiometric predictions from code created in the Thermal Synthesizer System (TSS). The thermal model outputs temperature profiles vs. time for numerous positions, and resulting gradients, which are then put into a Finite Element Structural model of the photometer. The structural model looks for thermally induced misalignments, including defocus, as a function of time. We care about these to the extent that they affect the star image on the focal plane, either in PSF shape or location. The output from the structural model is input to the Code V optical model of the *Kepler* optical system. As expected, because of our low expansion Graphite Cyanate construction and thermal design approach, thermal distortion effects were undetectable in both the shape and location of the PSF. By artificially increasing the amplitude of the thermal distortion effects, we determined that we could be sensitive to linear gradients across the Schmidt

corrector plate, and gradients across the focal plane mounting structure. Sensitivities in other areas (i.e. the spherical mirror) were smaller. These results will be used to generate specifications on the flight design for maximum allowable gradients.

A numerical model of the spacecraft, and the attitude control system were built. The attitude control system is capable of maintaining steady pointing over a specified bandwidth to a specified precision. Stars on the focal plane will show small movements with respect to the individual pixels of the CCD due to both jitter (higher frequency pointing errors) and drift (lower frequency pointing errors). The higher frequencies are due in part to the attitude control system control loop performance, and in part due to direct transmission of any spacecraft disturbances to the focal plane and optics. Based on previous Ball spacecraft, a model of the *Kepler* control system was built using the proposed attitude control hardware and the drift-away orbit environment. Although we use moderately small reaction wheels for *Kepler*, there is still some level of wheel "rumble" that is transmitted through the structure to the focal plane and optics. We used the previously described structural model, and the model's first 85 modes, to predict the level of disturbance at the focal plane that would result in apparent motion of a stellar image. We also evaluated the level of disturbance from the high gain antenna, which is moved once per day (1°) to track the Earth. Both the wheel disturbance and the antenna disturbance prove to introduce insignificant image motion (typically less than 1×10^{-8} radians per gram of input force at the source). The spacecraft attitude control system model is used to produce a pointing direction (LOS) vs. time curve for combined jitter and drift, to be used in the subsequent *Kepler* Performance model.

The photometer's numeric Performance Model is constructed in MATLAB. Inputs to this model include the predicted PSF at many locations across the focal plane, the spacecraft line of sight motion discussed above, inherent stellar variability, gain variations within and between pixels, dark current, pixel saturation, read noise, shutterless operation, etc. The Performance Model operates on a star field of 100 stars, with a range of brightness and relative position, and produces a realistic output signal vs. time for all pixels within a defined distance of the star. The output includes photon statistical variations (shot noise), read noise and dark current. The output is generated every 3 s for each pixel of each star region for 36 hours. The input starlight from some stars is decreased by an amount appropriate for a planetary transit across a star. The duration is programmable, typically for between 5 and 12 hours of the 36-hour stream of data. All pixel data streams are passed to a Fortran program that combines the data from all 100 stars into a synthetic frame for each exposure, adding readout smear. The synthetic frames are processed in the same manner as those in the laboratory demonstration. Thirty-six hours' worth of 15-minute flux data for all 100 stars is then passed to the transit extraction software (written by the author at NASA Ames), which calculates the transit statistics. The sensitivity of the transit detectability to variations in the performance model inputs can then be investigated. The planetary transits we have simulated show SNR's consistent with what we expected based on simpler RSS analysis of major error sources. They are also consistent with the laboratory results described earlier in this paper given the differences between the laboratory PSF and the flight PSF used in the numerical simulations.

6. ASSESSING TRANSIT DETECTABILITY

In this section we briefly describe the method used to assess detectability of transits in the numerical simulation. This treatment is similar to that presented in Koch et al. for examination of the laboratory data[2]. A detection statistic, *l*, in the simplest case of white noise and a deterministic signal is simply the dot product of the signal to be found with the data normalized by the noise variance where the time series are treated as vectors. More complex noise distributions require preprocessing before the dot product is taken to essentially whiten the noise. Hence, *l* is a real number that will be large and positive if the signal occurs in the data and the noise is small as the data vector and signal vector will be nearly parallel. Of equal interest is the value of *l* in the absence of the signal (the so-called null statistic). Transit detection basically is a problem of estimating the probability density distributions for these two cases. The degree of overlap of the two distributions determines the detectability of the signal.[8]

We injected Earth-sized, 5-hour transits into the flux time series resulting from the aperture photometry performed on the synthetic CCD frames and derived detection statistics after first removing a low-order polynomial trend from the data outside of each transit. Nine transits in all were injected sequentially into each star's light curve. The pixels for each star were examined prior to injecting the transits in order to account for the effects of saturation. Null statistics were calculated in the same manner but without adding the transits. Figure 7 presents the results for the baseline 5x5 edge PSF, demonstrating that the *Kepler* design meets the specifications for stars dimmer than 9[th] magnitude. The fact that the null statistic variances cluster near 1 is important, as it determines in large part the false alarm rate as a function of SNR. This result is similar to the laboratory simulation discussed at great length in Koch et al.[2] Figure 7b shows the corresponding results for the baseline run for the Kepler Testbed. Here, as in the numerical simulation, the single-event SNRs are adequate.
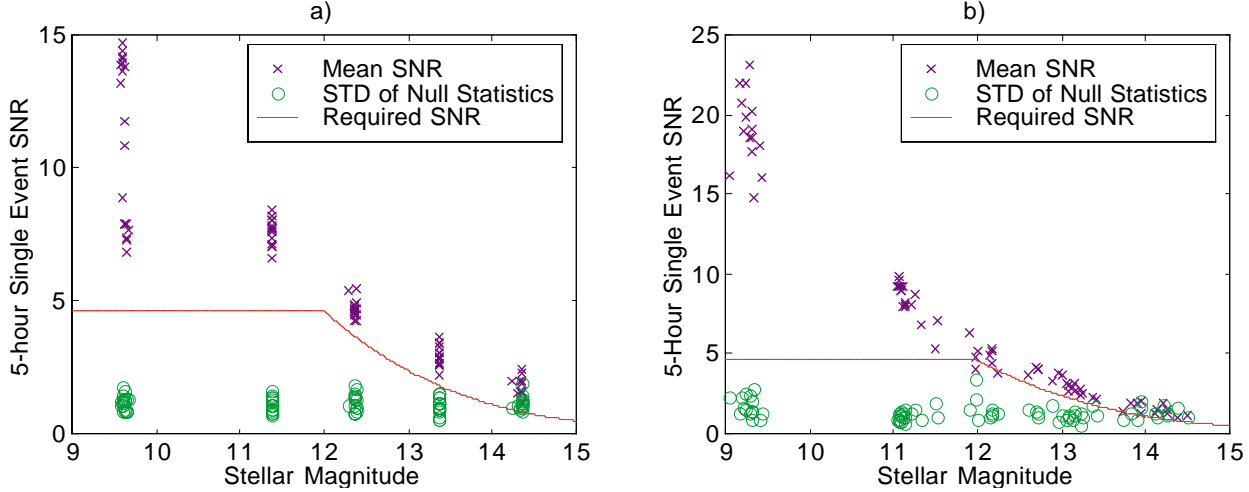
Figure 7. Single event 5-hour transit statistics for the baseline numerical simulation and for the baseline laboratory run. In panel a) the mean SNR for the numerical simulation is shown as X's for each star while the mean standard deviation of the null statistics (no transit present in data) is shown as O's, and the required SNR is shown by a solid line. In panel b) the corresponding statistics are shown for the baseline run in the laboratory Testbed (note the different vertical scale).

7. DATA COMPRESSION

We briefly explore the potential for lossless data compression techniques to allow transmission of star pixels with *Kepler's* data rate requirements. As a simple compression method, we took the first difference of the pixel time series for one of the lab test runs binned to 15 min and then examined the range of the excursions relative to the uncertainties. Figure 8 shows a histogram of the bits required to represent each pixel to a resolution of 1/4 the data uncertainties. Clearly, the average pixel requires only 5 bits at this level of resolution. For 170,000 target stars and 50 pixels per star, there would be 4 Gbits/day. This is well below the estimated 50 Gbits/day limit for 8 hours on a 34-m DSN station at *Kepler's* distance for the first two years. As the pixel time series are highly correlated, an algorithm such as a block wavelet transform followed by linear predictive coding and then entropy coding could result in substantial reductions in data volume.
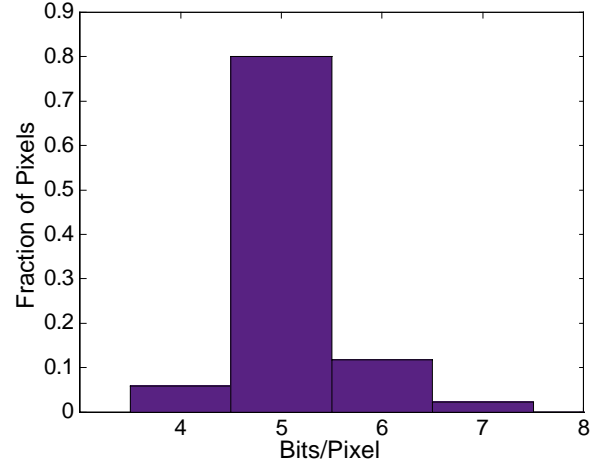

Figure 8. Bits per pixel required for lossless reconstruction of pixel brightness time series.

8. CONCLUSIONS

We have presented the theoretical underpinnings of the algorithms used in the *Kepler* Technical Demonstration to obtain 10 ppm differential aperture photometry on timescales of several hours with realistic optics, CCD operation, hardware-injected spacecraft motion and transits, and star field brightness distribution properties. This level of photometric performance is sufficient to detect stellar transits of Earth-sized planets with 1-meter class optics, such as proposed for the *Kepler* M*ission*.

The techniques described in this paper are applicable to photometry missions other than *Kepler*, and can result in significant improvements in data quality over simple aperture photometry. The calculation of stellar fluxes can be conducted either on-board the spacecraft or deferred to ground-based processing, depending on the balance between data link restrictions and onboard computing capabilities. For missions such as *Kepler*, where the FOV is the same for long periods of time, substantial reductions in data volume can be achieved by taking advantage of the redundancies in the data. We note that a method similar to the optimal pixel weighting scheme discussed in §3 is being applied to ground-based searches for 51-Peg-like planets by Tim Brown in the STellar Astrophysics & Research on Exoplanets (STARE) Project[9].

ACKNOWLEDGMENTS

REFERENCES

1. D. Koch, W. Borucki, L. Webster, E. Dunham, J. Jenkins, J. Marriott, and H. Reitsema, "*Kepler*: a space mission to detect earth-class exoplanets"*, SPIE Conference 3356, Space Telescopes and Instruments V*, pp. 599-607, 1998.
2. D. G. Koch, W. Borucki, E. Dunham, J. Jenkins, L. Webster, and F. Witteborn "CCD photometry tests for a mission to detect Earth-size planets in the extended solar neighborhood*", SPIE Conference 4013, Space Telescopes and Instruments VI*, 2000.
3. Robinson, L. B, M. Z. Wei, W. J. Borucki, E.W. Dunham, C. H. Ford, and A. F. Granados, "Test of CCD precision limits for differential photometry", *PASP* **107**, pp.1094-1098, 1995.
4. J.M. Jenkins, W.J. Borucki, E.W. Dunham, and J.S. McDonald, "High precision photometry with back-illuminated CCDs," *Proceedings of the Planets Beyond the Solar System and Next Generation Space Missions*, *A.S.P. Conference Series*. **119**, pp. 277-280, 1997.
5. M.H. Hayes, *Statistical Digital Signal Processing and Modeling,* John Wiley & Sons, New York, 1996.
6. S.B. Howell, "Two-dimensional aperture photometry - Signal-to-noise ratio of point-source observations and optimal data-extraction techniques", *PASP* **101**, pp. 616-622, 1989.
7. H.J. Deeg, and L.R. Doyle, "VAPHOT - a package for precision differential aperture photometry", *in Proceedings of the Photometric Extrasolar Planets Detection Workshop*, NASA Ames Research Center, in press, 2000.
8. R.N. McDonough, and A.D. Whalen, *Detection of Signals in Noise, 2nd Ed.*, Academic Press, San Diego, 1995.
9. T. M. Brown, personal communication, 1999.