

Data Analysis for the Periastron Passage of HD 80606b

Photometry and Ramp Correction

Our analysis used the Basic Calibrated Data provided by version S16.1.0 of the analysis pipeline at the Spitzer Science Center (SSC). We first applied the corrections recommended by the SSC to adjust for variations in pixel solid angle, and differences between flat-fielding of point sources versus extended sources. These corrections have only a very small impact on our analysis. Using the information contained in the FITS headers, we converted the measured fluxes to electrons, to facilitate the error analysis. We subtracted a background value for each 256x256 frame, based on a Gaussian fit to a histogram of pixel values. Typical background values were ~ 900 electrons per pixel, versus $\sim 400,000$ electrons of stellar signal, so fluctuations in the background do not dominate the photometric errors. Increases in pixel intensity due to energetic particle hits were more frequent in these 12-second exposures than for shorter exposures of brighter exoplanet stars. We corrected these pixels using a median filter applied to each pixel versus time. This was a two-stage correction, using a very large threshold at the first stage in order to correct the largest deviations, especially those that bloom over multiple pixels. The second stage threshold was set to 4σ , where the standard deviation was computed separately for each pixel versus time. A few frames had energetic particle hits that overlaid the stellar images to a degree that made their correction problematic. However, these frames are easily identified by their excursion in the final time series, and they were zero-weighted when binning the photometry over multiple frames.

Initial aperture photometry of each star used a square aperture of $(2N + 1) \times (2N + 1)$ pixels, including fractional pixels at the edges. We explored a range of aperture sizes (N), and we found that $N = 5$ gave the lowest noise photometry. The coordinates of aperture center were determined from a parabolic fit in each coordinate, using the brightest three

– 2 –

pixels in each coordinate for the fit. Figure S1 shows the photometric series for HD 80606 and HD 80607. This methodology is typical of high quality aperture photometry that we have done for other systems (e.g., Deming et al. 2007), and both the secondary eclipse of HD 80606, and its flash-heating near periastron, can already be discerned on Figure S1. Nevertheless we regard this initial photometry as inadequate for the present investigation. The intensities of each star increase by ~ 1 -percent throughout the time series. This is a well-known effect in Spitzer photometry at this wavelength, and is colloquially termed "the ramp". However, the ramps for HD 80606 and HD 80607 are distinctly different on Figure S1. In fact, the ramp for HD 80607 does not increase monotonically, and such behavior is not normally encountered in exoplanet photometry at this wavelength. The different ramp shapes for the two stars preclude using HD 80607 to correct the ramp in HD 80606.

We attribute the different ramp shapes to the varying positions of the stars relative to the boundaries of their surrounding pixels. This can produce different ramp shapes in each pixel depending on the illumination history of each specific pixel, i.e., whether it is receiving an increasing or decreasing photon flux with time as the image drifts in position. Hence the integration over pixels inherent in aperture photometry can produce different baselines in aperture photometry for stars at different positions. An example suffices to illustrate this point. Suppose that one star is initially centered on a given pixel. Then the detector ramp (essentially, the gain) for that pixel will reach its full value relatively quickly. In the contrasting case where the other star is centered mid-way between two pixels, those pixels will reach full sensitivity relatively slowly. Hence, aperture photometry in those two cases will produce different ramp shapes. To fully account for such effects, we have to correct the ramp separately for each pixel, taking into account the illumination history of each pixel.

Another deficiency of our initial photometry is that the positions of the two stars derived from the parabolic fit did not vary in unison as perfectly as we expected. Hence we

– 3 –

undertook a more sophisticated data analysis for our final results.

The final photometric analysis removed the detector ramp separately for each detector pixel, and we used better image centroid positions for both stars. We adopted the IRAC point spread functions convolved with the pixel response, available as PRF FITS files on the SSC website. We resampled these PRFs to much higher spatial resolution, shifted them in position to match a candidate stellar position, rebinned them back to the observed pixel resolution, and fit them to the data. We repeated this process using a Levenberg-Marquardt algorithm to find the X,Y stellar position having the minimum χ^2 . We found that the PRFs provided by the SSC do not give particularly good fits to our data. We believe that the differences are due to electronic pixel-reading effects such as the "bandwidth effect" described in the IRAC Data Manual. We therefore constructed custom PRF functions. We re-mapped the differences in the fit for each image back to the PRF, and averaged these corrections for all 8051 images to yield improved PRF functions for each star. A Levenberg-Marquardt re-fit of these improved PRFs gave differences comparable to the photon noise, and also gave image centroid positions that varied closely in unison between the two stars.

Figure S2 shows the Y-position of each star versus time from the PRF fitting, and Figure S3 shows the X-position, which varies much less. An oscillatory behavior with a ~ 1 -hour period is superposed on a slow monotonic drift amounting to ~ 0.5 -pixels in Y. This oscillation has been seen in many previous Spitzer investigations, and is due to telescope pointing jitter. The slow, monotonic drift in position was also seen by Knutson et al. (2007), but its origin is not understood. Regardless of its origin, it has implications for the ramp correction. We correct the detector ramps on a per-pixel basis, i.e. we fit a ramp function to the time series of each pixel intensity. Since the ramp is essentially a gain variation, we correct the measured signal for image motion prior to fitting the ramp

– 4 –

function.

Let the total photon flux (star+planet) incident *on a given pixel* at time t be denoted as $P(t)$. $P(t)$ varies primarily due to motion of the image relative to this specific pixel. Let the signal produced by the pixel be $S(t) = S_s(t) + S_p(t)$, where s indexes the portion of the signal from the star, and p indexes the portion from the planet. Spitzer $8\ \mu\text{m}$ photometry is known to be affected by the detector ramp, that behaves like a time dependent gain, $G(t)$ (Deming et al. 2006, Knutson et al. 2007, Harrington et al. 2007, Deming et al. 2007, Demory et al. 2007). The signal produced by the pixel is then given as:

$$S(t) = P(t)G(t) \quad (1)$$

Because the centroid position of the stellar image is well-determined by the PSF-fitting procedure described above, we use that information to determine variations in $P(t)$ to first order. We write:

$$P(t) = 1 + a_1\delta Y + a_2\delta X, \quad (2)$$

where $\delta Y = Y - Y_{av}$ is the difference in Y-coordinate of the image centroid from its average value (Y_{av}), and similarly for X . The 1 on the right side of Eq. (2) results from normalizing the average total signal for that pixel to unity. The $a_1\delta Y$ and $a_2\delta X$ terms are significantly less than unity, because δY and δX are a small fraction of the width of the Spitzer PSF at this wavelength. When making the first order determination of $P(t)$, we use only a low frequency version (see below) of the Y-term, since the image motion is overwhelmingly in the Y-direction (Figures S2 and S3). Note that there is physical flux conservation, i.e. the sum of the photons incident on all pixels is independent of image motion per se. We enforce flux conservation to first order in the analysis by deriving a_1 from a centered finite-difference

– 5 –

equation for the derivative of the image intensity with respect to Y . The difference equation has the property that the sum of the a_1 values over all pixels is zero, so modeled flux is conserved as it is transferred from pixel to pixel by Y -direction image motion.

The properties of $G(t)$ (the ramp) are well established from several Spitzer exoplanet photometric investigations. It is often modeled as a sum of a linear and a logarithmic term:

$$G(t) = 1 + a_3t + a_4\ln(t) \quad (3)$$

Therefore, neglecting the planet signal for the moment, the expression for $S_s(t)$ becomes:

$$S_s(t) = 1 + a_1\delta Y + a_2\delta X + a_3t + a_4\ln(t) \quad (4)$$

where we have dropped higher order terms in the cross multiplication (like $a_1a_3t\delta Y$) for the time being (see discussion below).

We now explicitly include a term for the heating of the planet. We also divide the X - and Y - dependent terms into low and high temporal frequency components in order to allow for a better fit to the data. (The low frequency position drift and the high frequency position jitter do not vary in unison, and must be accounted for separately in the analysis.) To separate these components, we fit a fourth order polynomial to the X - and Y - coordinates versus time, and we use that polynomial to represent the low frequency drift in time (see Figures S2 and S3). We experimented with the order of the polynomial fit, and settled on fourth order as representing the lowest order that adequately represents the gradual drift in the coordinate values. Subtracting the gradual changes represented by the polynomial, gives the high frequency position jitter that is dominated by the 1-hour telescope oscillation. Including these terms, the signal equation for a given pixel becomes:

– 6 –

$$S(t) = 1 + a_1\delta Y_l + a_2\delta X_l + a_3t + a_4\ln(t) + a_5\delta Y_h + a_6\delta X_h + a_7\gamma(t), \quad (5)$$

where $\gamma(t)$ is the function that describes the time dependence of the planet's heating. The indices l and h index the low and high frequency components of image motion, as described above. The solution for the planet signal is made by solving the above equation for the best-fit a_i coefficients using multivariate linear regression on the data for 8050 values of time (t) at each pixel, except that the $a_1\delta Y_l$ term is calculated *a priori* under the constraint of flux conservation when summing all pixels. There are several reasons for calculating this term *a priori*. First, the low frequency drift in Y is closely linear in time (see Figure S2), and we found a tendency for $a_1\delta Y_l$ to be degenerate with a_3t in the regression solution. By calculating $a_1\delta Y_l$ *a priori*, we remove that degeneracy, and we can also impose flux conservation on the solutions, at least to first order. Solving for the remaining terms is done by linear regression, independently for each pixel. The regression fit calls the IDL routine REGRESS, which utilizes matrix inversion. While independent regressions for each pixel formally allow flux conservation to be violated (e.g., by the h -subscripted terms), the flux variations implied by those remaining terms are not large in comparison to the noise level. Moreover, the extracted planet signal is independent of these variations, since our planet result is expressed as a ratio to the stellar (i.e., baseline) signal (flux variations are common to both star and planet signal).

We explored the effect of including the higher order cross terms, like $a_1a_3t\delta Y_h$, in the solution of Eq. (5). Since these terms are a small perturbation to the initial solution, we included them using an iterative approach. Solving Eq. (5) for the a_i , we calculated the magnitude of the missing higher order terms. Adding those to the right side of Eq. (5), is equivalent to subtracting them from the left. We subtracted a fraction of the missing terms from the left hand side of Eq. (5), where the fraction (typically, half) was adjusted

– 7 –

to facilitate convergence of the iteration without overshoot. We re-solved for the new a_i , re-computed the new higher order coefficients, and re-evaluated Eq. (5) with their full values included. We continue this process iteratively until the relative error in $S(t)$ for that pixel is negligible (much less than a single electron count) at all times. We also explored including the higher order terms directly in the linear regressions, although that does not constrain their coefficients to equal the product of the corresponding coefficients from the first order terms. However, we found that neither approach produced significant reduction in the χ^2 of the fit, so inclusion of the higher order terms is not useful. We attribute this to imprecision (i.e., noise) in the independent variables, for example the X- and Y-coordinates of the image (Figures S2 and S3). The $a_1 a_3 t \delta Y_h$ term, for example, quantifies effects related to the ramp-induced gain increase on the intensity perturbations resulting from Y-axis pointing jitter, i.e. a gain perturbation to an intensity perturbation. To include such a small effect precisely, we need to know the independent variable (i.e., Y-amplitude of the pointing jitter, (δY_h)) very precisely. But it is clear from Figure S2 that measuring the amplitude of this pointing jitter at a specific time t is imperfect due to measurement noise. Hence we concluded that including higher orders terms in Eq. (5) is neither possible nor desirable for our analysis.

After solving the first order signal equation (Eq. 5) for each pixel, we constructed a series of "ramp images" wherein each pixel was assigned the value of the fitted ramp at that time, i.e. the right side of Eq. (5) *with the planet term omitted*, and properly re-normalized relative to other pixels. We then performed aperture photometry on the ramp images, as well as on the data before decomposition of the signal by Eq. (5). This yields two time series for each star: a time series of the data including the ramp and the planet, and a time series of the ramp function obtained from solving Eq. (5). By dividing the former by the latter we isolate the planet signal and cancel the ramp. This procedure produced the results illustrated in the main paper, and we note that it produces a flat time series for HD 80607,

– 8 –

which functions as a control case.

In solving Eq. (5), it is necessary to have an initial estimate for the $\gamma(t)$ function that represents the time dependence of the planet signal. Note that it is impossible to solve for the detector ramp and the planet signal independently without making *some* assumptions concerning their functional form. Fortunately, the detector ramp is sufficiently well-studied by other investigations that we can be confident of modeling it by the linear plus logarithmic terms in Eq. (5). Also, because we designed the observations so that periastron occurred near the end of the time series, the detector ramp at this point was relatively flat. Therefore we can make an initial estimate for $\gamma(t)$, *solely from the data*, without recourse to model calculations. We derived this initial estimate by fitting a linear plus logarithmic baseline to the simple aperture photometry (Figure S1), and deriving $\gamma(t)$ from the deviations from this fit. The quasi-discontinuity that is already apparent near $JD = 2454424.8$ in Figure S1 represents the effect of the planet, and a fit using a nominal ramp shape gives us an initial estimate for $\gamma(t)$, independent of any model.

Beginning with a solely data-based estimate for the $\gamma(t)$ term is important, since otherwise the analysis could potentially produce false results by adjusting the coefficients of the non-planet terms to accommodate a predicted planet signal from the model. Using our initial data-based estimate for $\gamma(t)$, we did our complete analysis as described above, and extracted the planet signal with greater fidelity. This signal turns out to have a somewhat different shape than the initial estimate for $\gamma(t)$. We then fit a model to the derived planet signal, and we replace the initial estimate for $\gamma(t)$ by the model fit (as a revised $\gamma(t)$), and re-do the entire analysis in an iterative fashion. Using the model fit in this manner is necessary, because only a smooth function will ensure that noise in the planet signal does not propagate and amplify. Recall that we derive the planet signal as the ratio of two photometric time series (described above). The numerator in this ratio is aperture

– 9 –

photometry on the baseline-uncorrected data. The denominator is aperture photometry on the baseline, i.e. the right side of Equation (5), *with the planet term removed*. Thus, inclusion of $\gamma(t)$ in the regressions allows the analysis to gracefully account for the presence of the planet, and prevents the derived baseline from being distorted by it. We did a number of tests to insure that our final result was not sensitive to usage of a model to represent $\gamma(t)$. We perturbed the model-based shape of $\gamma(t)$ in three ways. First we added a sine wave of amplitude 0.01-percent, with a 1000 data frame period and random phase. Inclusion of this ripple on $\gamma(t)$ produced negligible effect on our final results. Two other tests involved distorting $\gamma(t)$ by squaring it, and also using its square-root. These distortions also had negligible effect on our final result.

We experimented with other variants on the linear regressions. For example, we tried using individual X- and Y-pixels positions for the two stars, versus the default method of averaging their position changes. We also evaluated the inclusion of a logarithmic-squared term in time, but it did not improve the results significantly. We evaluated variants on the fitting procedure based on the scatter in the HD 80607 results (since it has no planet to contribute to the scatter), but also by relying on our subjective judgment of how well each fit accounted for the known systematic effects such as the ramp and the telescope oscillation.

Tables S1 and S2 list the individual photometry values for both HD 80606 and HD 80607 as well the weights used in binning the data for the Figures. The weights were taken to be proportional to the inverse of the calculated variance of each point from the statistical error model (see below). A few outlying points, affected by uncorrectable energetic particle hits, were given zero weight.

Random Error Analysis

We developed a statistical error model for the observations. In each observed frame

– 10 –

we calculate the predicted photometric error of the aperture photometry, based on the Poisson fluctuations in the number of detected electrons (dominant), as well as read noise and background fluctuations (smaller). As a test, we apply the model to HD 80607, since it is not expected to have any real fluctuations that might masquerade as observational error. Figure S4 shows the distribution of photometric fluctuations for all 8051 observations of HD 80607, normalized to the predicted error of each observation. We fit a Gaussian to this distribution and find an excellent fit using a standard deviation of 1.25, i.e., the photometry has a scatter 25-percent greater than predicted from the error model. This is excellent performance, and is consistent with other Spitzer exoplanet photometry at this wavelength.

The heating of the planet at periastron, and the secondary eclipse, are obvious on Figure 2 of the main manuscript. Nevertheless, we here go beyond a ‘chi-by-eye’ analysis, and we discuss quantitatively the statistical reality of the planet’s heating (the reality of the eclipse is discussed in a separate section below). The error bars for the binned data shown in Figure 2 were calculated from the internal scatter of the observations within each bin. As noted above, they approach the theoretical limiting errors from photon statistics. HD 80607, having no heating, is particularly well-suited to evaluate whether the errors decrease with the square-root of time, as expected for white noise. A key test is whether the error bars - calculated from the internal scatter within each bin - are consistent with the actual scatter of the binned points. We find that the actual scatter is 16% greater than the average error bar for HD 80607. We regard this as outstanding consistency, but it may indicate the existence of a small amount of ‘red noise’. We evaluate the statistical significance of the planet’s heating by conservatively increasing the error bars for HD 80606 on Figure 2 by 16%. Excluding the data with the planet in eclipse, we evaluate the null hypothesis of no heating by calculating the chi-squared for a constant flux model using the binned data. This yields a chi-squared of 536.9 for 353 degrees of freedom, rejecting the null hypothesis at a 6σ level of significance. Moreover, this actually understates the

– 11 –

significance of the observed flash-heating, since it only implies that there is heating of the planet at some time during the observed period: it does not factor-in that the observed heating occurs in close coincidence with the known time of periastron passage.

Limits of the Analysis

Our analysis is designed to correct for the known systematic error in Spitzer photometry at this wavelength, namely the ramp-up in detector sensitivity as a function of time. Our regression method is similar to a high-pass numerical filter, wherein the eclipse of the planet (see below), and the rapid heating near periastron, are separated from the detector ramp. The ramp has a longer time scale, and a very different functional dependence on time, and can be separated on this basis. Although our analysis is very successful in making this separation, some aspects of our results are more uncertain than others. The ramp increases most strongly at early times, and the planet signal is increasing slowly then. So our results at early times are more uncertain than our results near periastron, which are quite robust.

Two aspects of the "early planet" need specific discussion. First, a slight decrease in planet emission, not predicted by the model, is apparent in the time span before eclipse. If this is real, it might indicate (for example) that the heated regions of the planet were turning away from us due to rotation. But linear least squares reveals that the rate of decrease is $-9.3 \pm 3.1 \times 10^{-4}$, i.e. a 3σ effect. Given that the early data are the most difficult part of our analysis, we do not claim that this apparent decrease is real. A second aspect is the fact that the early planet level (leftmost points on Figure 2 of the main paper) is greater than the "no planet" level (the in-eclipse level on Figure 2 of the main paper). We believe this is a real effect, and is due to the tidal luminosity of the planet, that causes it to have an appreciable temperature even under low stellar irradiation. The tests wherein we distorted $\gamma(t)$ (see above) all produced similar results where the early planet emission exceeded the no-planet level. Those tests are the most applicable way to probe the reality

– 12 –

of this phenomenon. Hence we believe that we are observing the tidal luminosity of the planet. But because this conclusion derives from the most difficult part of our analysis, we conservatively claim only that our result is consistent with theoretical predictions for the temperature induced by the planet's internal tidal dissipation ($T \sim 700K$, Wu & Murray 2003).

The Secondary Eclipse

Figure S5 shows the secondary eclipse. As an initial test for the reality of the eclipse, we did bootstrap trials wherein we permuted the unbinned deviations from the model fit to make 50,000 sets of fake data. For each set of fake data, we search for eclipses using a boxcar fit as a function of central eclipse time, varying the eclipse time over the entire 30-hour range of the data. We tabulate the maximum eclipse depth found by this algorithm for each fake data set. In the 50,000 trials, not a single instance could be found where the boxcar fit derived an amplitude as large as observed in the real data. This is especially significant because the boxcar algorithm was given the freedom to vary the central eclipse time, whereas the eclipse in the real data occurs very close to the time predicted by fits to the radial velocity data. We conclude that the reality of the secondary eclipse in our photometry is unequivocal.

To derive the central eclipse time, duration, depth, and ingress/egress time, we fit a set of synthetic eclipse curves that span a 3-dimensional grid of input parameters. These parameters are central time, duration, and ingress/egress time. The eclipse shape was synthesized at each grid point by connecting the contact points with straight lines, and assuming a constant transverse velocity for the planet. The high eccentricity of the orbit causes a variation in transverse velocity during eclipse, translating to a difference in ingress versus egress time (whereas we assume they are the same). However, the error in our method is negligible, considering the precision of the eclipse detection ($\sim 7\sigma$), and the

– 13 –

fact that the random error on the ingress/egress time slightly exceeds the ingress time itself (see below). At each grid point (i.e., values of central time, eclipse duration, and ingress/egress time) we find the best-fit eclipse depth using linear least squares to scale the synthesized eclipse curve to variable depth, adopting a baseline from the model fit to the out-of-eclipse data (Figure 2 in the main paper). We allow the fit to apply a scale factor to the model baseline to account for the observed fact that the planet apparently heats faster than the model predicts near eclipse. In doing the least-squares fits for eclipse depth, we always use unbinned data. But the eclipse is easier to see "by eye" in binned data, as illustrated in Figure S5. We also solved for the eclipse parameters using the data prior to ramp correction, and derived the baseline via a simple parabola fit via linear regression. This check on our eclipse analysis produced the same results, within the errors.

Given the level of observational error, it was possible to choose the grid spacing so that the entire grid could be sampled in a relatively rapid computation, and we estimated errors by repeating the calculation for 10,000 fake data sets having the same precision. On this basis we find the eclipse central time, ingress/egress time and duration: $T_c = 2454424.736 \pm 0.0038$, ingress/egress time = 0.12 ± 0.12 hours, eclipse duration = 1.68 ± 0.24 hours, and we find the depth of the eclipse = 0.00130 ± 0.00018 . The best-fit slope of the baseline was 2.3 times the prediction from the model, indicating a more rapid apparent heating than anticipated. Scaling the flux from a 5750/4.5 Kurucz model atmosphere to the $T_{eff} = 5645$ listed by Naef et al. (2001) for the star, we calculate that this eclipse depth corresponds to a planet brightness temperature of $T_b = 1300 \pm 170K$, assuming solar and Jupiter values for the stellar and planetary radii, and where the error reflects only the precision of the observed eclipse depth.

A Non-Hydrodynamical Model of the Planetary Radiative Response

The primary effect that we observe is the $8\mu\text{m}$ flux increase from the planet during the

– 14 –

time interval surrounding periastron passage. We model this response with a joint radiative and hydrodynamical model, with the results shown in Figure 2 of the main article. Because the photometric observations ended soon after periastron, there was insufficient time to observe the full hydrodynamic response of the planet, and it is instructive to compare the photometry with models in which only the global radiative response is modeled.

Figure S6 illustrates two model responses in which the dynamics have been turned off. In the first and most extreme case (shown in red) we assumed instantaneous thermal re-radiation of the absorbed stellar flux. This approximation does a poor job of fitting the data, confirming that the atmosphere indeed contains a non-zero thermal inertia. In the second case (shown in dark blue) we use our present radiative model. For this case, the fit is comparable to the fit obtained with the full radiation-hydrodynamic model, indicating that Spitzer is primarily observing a radiative, rather than hydrodynamic response from the planet.

– 15 –

REFERENCES

- Deming, D., Harrington, J., Seager, S., & Richardson, L. J. 2007, *ApJ*, 667, L199.
- Deming, D., Harrington, J., Laughlin, G., Seager, S., Navarro, S. B., Bowman, W. C., & Horning, K. 2007, *ApJ*, 667, L199.
- Demory, B.-O., & 13 co-authors, 2007, *A&A*, 475, 1125.
- Harrington, J., Luszcz, S., Seager, S., Deming, D., & Richardson, L. J. 2007, *Nature*, 447, 691.
- Knutson, H. A., Charbonneau, D., Allen, L. E., Fortney, J. J., Agol, E., Cowan, N. B., Showman, A. P., Cooper, C. S., & Megeath, S. T. 2007, *Nature*, 447, 183.
- Langton, J., & Laughlin, G. 2008, arXiv:0808.3118.
- Naef, D., & 11 Co-authors, 2001, *A&A*, 375, L27.
- Wu, Y., & Murray, N. 2003, *ApJ*, 589, 605.

This manuscript was prepared with the AAS L^AT_EX macros v5.2.

– 16 –




Fig. 1.— Spitzer photometry of HD 80606 and HD 80607 without correction of the ramp. Note that the secondary eclipse and the heating of the planet near $\text{JD}=2454424.7$ are apparent even prior to the removal of the detector ramp.

- 17 -




Fig. 2.— Y-pixel positions of HD 80606 and HD 80607 versus time. Note the oscillatory behavior due to an oscillation in telescope pointing with an approximate 1-hour period, as well as a long-term drift amounting to ~ 0.5 -pixels. The absolute values of the Y-coordinates plotted here are relative to a subarray, and do not refer to the original data coordinates. The red line is the fit of a 4th order polynomial, intended to isolate the gradual component of image motion.

– 18 –




Fig. 3.— Same as Figure 2, but for the X-coordinate of image position.

– 19 –




Fig. 4.— Histogram (points) of the deviations from the average value for the control star, HD 80607, tabulated in units of the predicted error. The curve is the fit of a Gaussian, and the best-fit retrieves a standard deviation for the distribution = 1.25, i.e., scatter in excess of the predicted errors by 25-percent.

– 20 –




Fig. 5.— Final photometry for HD 80606, binned in groups of 50 points, showing the secondary eclipse, centered at $HJD = 2454424.736$. The curve is the fitted eclipse as described in the text. The slope in the baseline is due to the rapid heating of the planet as it approaches periastron.

– 21 –




Fig. 6.— Comparison of the HD 80606 photometry with purely radiative models for the planetary response. The red curve corresponds to a model that assumes instantaneous re-radiation of the absorbed stellar flux. The blue curve corresponds to our present radiative model (as described in the text and in Langton & Laughlin 2008). Time is measured in HJD, relative to HJD 2454400.