Aligning the ZTF science focal plane using stellar images

Gina E. Duggan, Richard Dekany, Jennifer Milburn

Gina E. Duggan, Richard Dekany, Jennifer Milburn, "Aligning the ZTF science focal plane using stellar images," Proc. SPIE 10702, Ground-based and Airborne Instrumentation for Astronomy VII, 107024K (6 July 2018); doi: 10.1117/12.2309994

SPIE.

Event: SPIE Astronomical Telescopes + Instrumentation, 2018, Austin, Texas, United States

Aligning the ZTF Science Focal Plane Using Stellar Images

Gina E. Duggan, Richard Dekany and Jennifer Milburn

Caltech Optical Observatories
California Institute of Technology
Pasadena, CA 91125 USA

ABSTRACT

The Zwicky Transient Facility (ZTF) is a next-generation, optical, synoptic survey that leverages the success of the Palomar Transient Factory (PTF). ZTF has a large science focal plane (SFP) that needs to be aligned such that all portions of the CCDs are simultaneously placed in focus to optimize the survey's efficiency. The SFP consists of 16 large, wafer-scale science CCDs, which are mosaicked to achieve 47 $\deg^2$ field of view. The SFP is aligned by repositioning each CCD based on the measured height map, which is a map of the camera's z position at which each portion of the CCD is in focus. This height map is measured using on-sky stellar images in order to recreate the optical path that will be used throughout the survey. We present our technique for placing the SFP in focus, which includes two different methods to measure the height map of the SFP. The first method measures the height at which a star is in focus by fitting a parabola to each star's photometric width as the star is moved in and out of focus. The second method measures the height by decomposing a defocused star into its image moments. We will discuss the strengths and limitations of each method and their outputs. By repositioning the CCDs, we were able to reduce the standard deviation of the height map from 33 to 14 microns, which improved the survey's speed by $\sim 81\%$.

Keywords: ground-based telescopes, instrumentation, ccd cameras, ccd positioning, ccd alignment, metrology

1. INTRODUCTION

ZTF is a high-efficiency, wide-field, optical, time-domain survey developed to explore new parameter space for short-cadence science and to develop scientific utilization infrastructure in preparation for the Large Synoptic Survey Telescope (LSST). The ZTF Observing System (OS) consists of the 1.2-meter diameter Samuel Oschin Telescope at Palomar Mountain. The OS is outfitted with a ZTF Camera[1] containing a new 640 Mpix 16 CCD mosaic science focal plane (SFP). The SFP is housed in a vacuum cryostat, positioned with a hexapod, and located at the prime focus of the Schmidt telescope system.

ZTF is an updated version of the Palomar Transient Factory (PTF) with new hardware and software. ZTF is designed to dramatically improve the survey speed of PTF. Improvements, including new CCDs and updated telescope slewing mechanisms, enable the exposure time per image and overhead between images to be decreased by a factor of two and three, respectively. Arguably the most important improvement is increasing the field of view by a factor of seven—from 7.26 to 47 $\deg^2$. All of these improvements combine to provide an areal survey rate of 3,750 $\deg^2$/hr, which will enable ZTF to scan the viewable sky (3π) an average of 290 times per year.

The 47-$\deg^2$ SFP presents a technical challenge: positioning all areas of the SFP in focus simultaneously. This is important for the survey's sensitivity, since having a portion of the SFP always out of focus reduces the average limiting magnitude of the stars ZTF is able to observe or increases the time needed to achieve a given limiting magnitude. The engineering requirement is to measure the SFP with 10 μm precision in order to achieve the desired photometric sensitivity.

We present two different methods to create a height map, which is a map proportional to the height of each pixel across the entire SFP. Both of these methods utilize on-sky images. This provides a distinct advantage over lab profilometry imaging, because we are replicating the conditions that will be used throughout the survey.

Further author information: (Send correspondence to G.E.D.)
G.E.D.: E-mail: gduggan@astro.caltech.edu

Ground-based and Airborne Instrumentation for Astronomy VII, edited by Christopher J. Evans, Luc Simard, Hideki Takami,
Proc. of SPIE Vol. 10702, 107024K · © 2018 SPIE · CCC code: 0277-786X/18/$18 · doi: 10.1117/12.2309994

The final SFP height map is an average of the maps produced by each method. The ZTF camera was brought back to the lab to re-shim the CCD supports to minimize height differences across the SFP. Then the ZTF camera was re-installed, and we report the improvement to the height variation across the SFP and the resulting improvement in survey efficiency.

2. PLACING THE SCIENCE FOCAL PLANE IN APPROXIMATE FOCUS

In May 2017, the ZTF cryostat was installed into the Oschin Telescope in a temporary optical configuration, which did not fully conform to the ZTF OS optical design.[2] Due to vendor delays, a large aspheric plate, needed to compensate for spherical aberration induced by the thick cryostat window, was not available for these tests. Based on Zemax design analysis, the absence of this aspheric plate resulted in aberrations present at all points of the ZTF optical field that are described by a Seidel spherical wavefront aberration coefficient of $W_{040} \sim -7.0$ waves. With a R-band filter, this results in a shift of focus to achieve the minimum root mean square (RMS) spot size of

$$\delta z = 8(F/\#)^2 \times \left(-\frac{4}{3}W_{040}\right) = -\frac{32}{3}(F/\#)^2 W_{040} = -\frac{32}{3}(2.5)^2(-7 \times 0.6\mu\text{m}) = 280\mu\text{m}, \quad (1)$$

relative to the paraxial focus that would be optimal in the absence of any optical aberrations. Here, we define positive motion to be in the direction from the prime focus toward the telescope primary mirror. However, a global shift in focus does not effect our measurements at all because we are measuring the relative heights in the SFP. The hexapod height at which each part of a CCD is in focus relative to one another (described by a 'height map') is independent of the large spherical aberration present during these experiments. The combination of this large spherical aberration and vignetting does effect our measurements, and we discuss this further in Secs. 3.1.2 and 3.2.2.

After installing the positioning hexapod onto the prime focus hub, the cryostat onto the hexapod, and the R-band optical filter to the front of the cryostat, we performed a coarse, initial, solid-body alignment of the cryostat relative to the telescope focal plane. We manually adjusted the hexapod z*, tip, and tilt using field- and hexapod-z-dependent stellar image widths as the metric—see this surface map of the stellar full width at half maximum (FWHM) given in Fig. 1 as an example. By minimizing the stellar FWHMs seen in these maps, we are able to distinguish the optimal hexapod z position (i.e., the z position that places the majority of the SFP in focus) within 35 μm in variable seeing. Although roughly identifying the optimal z position is important, adjusting the tip and tilt to flatten the SFP is essential to measuring an accurate height map of the SFP. A visible gradient on the stellar FWHM map corresponds to a large global tip and/or tilt, which dramatically increases the range of z positions needed to place each section of the SFP in focus.

Following this, we obtained several series of through-focus scans by taking images at different z positions, which included positions above and below the optimal z position for the SFP that was found using our stellar FWHM map. We varied the camera exposure time of these images between 30 and 120 sec and varied the number of the different z positions between five and nine positions. This was done to explore both the dependency of our prepared analysis methods to out-of-focus distance and signal-to-noise ratio (S/N), which is a function of the number of sky- and read-noise pixels within the defocused stellar point spread function (PSF). We will discuss the optimal set of images selected for each method in the following sections.

3. CREATING A HEIGHT MAP OF THE SCIENCE FOCAL PLANE

The ultimate goal is to calculate a set of height adjustments to the CCDs required to bring the entire SFP into focus simultaneously. This is equivalent to calculating the hexapod z position (or height) where each section of the SFP is placed in focus. For example, if a section of the SFP reaches focus at a larger z value (relative to the rest of the SFP), one has to move that section of the CCD up relative to the cryostat body (or e.g. the cryostat cold plate) to flatten the plane. In other words, the 'height' in height map refers to the height of the hexapod

*Typically we refer to the hexapod axial as the 'focus'. However, to avoid confusion with the state of a star being 'in focus' and the hexapod position where that occurs, we will simply refer to this axis as z.

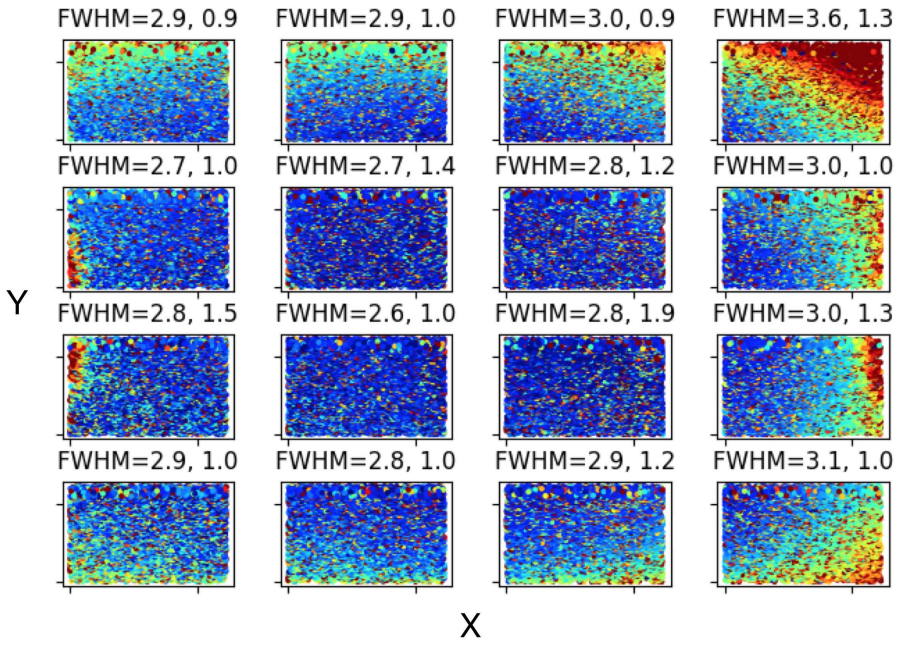

Figure 1. Stellar FWHM map of the SFP. This field-dependent stellar image FWHM was measured at the optimal hexapod z position obtained using rigid-body motions of the ZTF hexapod. This technique is used to place the SFP in approximate focus by adjusting the z, tip, and tilt of the hexapod. The FWHM color ranges from dark blue (FWHM = 2 pixels) to dark red (FWHM = 4 pixels). The FWHM average and standard deviation for each CCD is printed above each CCD (in the form FWHM = mean, std). With this method we estimate that we can identify the z position that places the majority of the SFP in focus within 35 μm.

when the section is in focus and is actually the inverted map of the true shape of the CCDs, so the height map is equivalent to the adjustments required to flatten the SFP. This section discusses how we create this height map of the SFP.

We generalize "each section" of the SFP to be each N x N pixel sub-frame of the image. In practice N=512 was found to be the ideal sub-frame size, since it assures that a sufficient number of stars are present in the determination of the height for that section. The end result is a 48 x 48 pixel image with each CCD represented by a 12x12 sub-array and a total of 2,304 sections in the SFP.

Although all the pre- and post-processing is identical, we can determine the height for each section of the SFP by two different methods: parabola method and donut method. Each method is impacted differently by the sources of biases present in ZTF. Ultimately, these methods were in agreement and one did not appear inherently less biased, so we simply averaged the results of the two methods. We will discuss the theory behind each method, the main sources of bias, and the detailed procedure used to calculate a height map.

3.1 Parabola method: parabolic fits to FWHM

With the parabola method, the height at an individual star's position is determined by fitting a parabola to FWHM measurements of the star in images at different hexapod z positions. We tested this method in October 2015 on PTF to confirm the method's ability to measure a height map within ZTF's required precision of 10 μm.

3.1.1 Parabola algorithm

In the absence of diffraction and atmospheric or static design optical aberrations, we would expect the dependence of the stellar FWHM on z, as shown in Fig. 2, to be symmetric on both sides of focus (i.e., the hexapod z that places the star in focus, z_0). The aforementioned diffraction and aberrations, however, preclude the experimental reproduction of the idealized curve, and add a 'floor' value for RMS spot size ($s_{\mathrm{rms,min}}$). The RMS image radius,

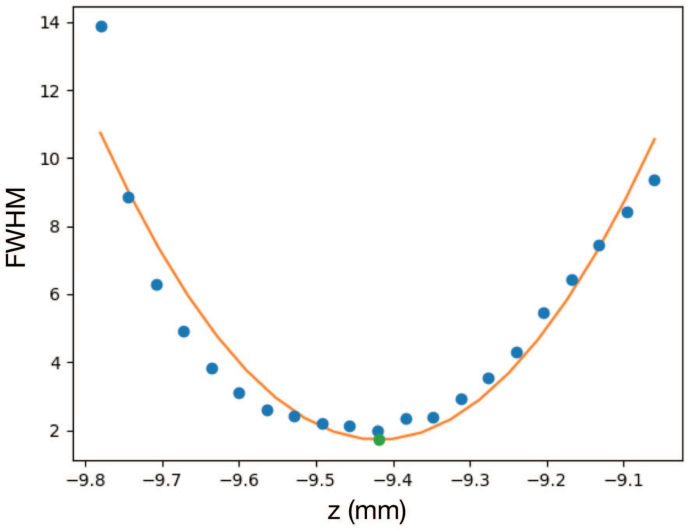

Figure 2. Measuring the height for a ZTF star using the parabola method. Here we fit a parabola (orange) to a single star's FWHM measured in many images (blue dots) to calculate the height where the hexapod places the star in focus (z_0, green dot). By averaging the standard deviation of the height returned by the parabolic fit for all stars, we measure an average height error of 15 μm for a single star. The shape deviation from an ideal parabola is due to the combination of spherical aberrations and vignetting.

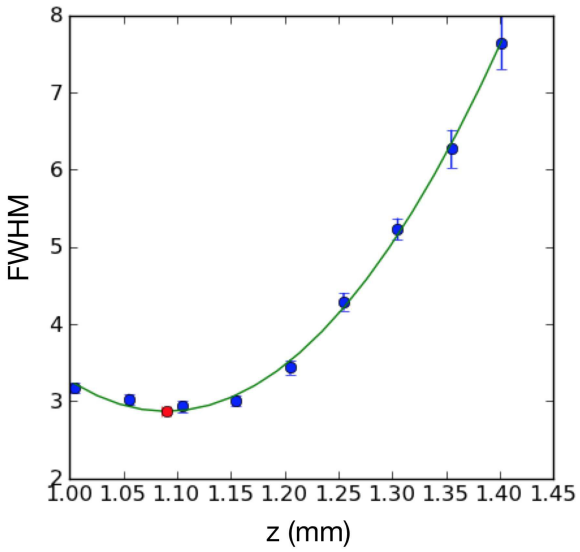

Figure 3. Measuring the height for a PTF star using the parabola method. We fit a parabola (green) to a single star's FWHM vs. z (blue dots) to calculate the height where the hexapod places the star in focus (z_0, red dot). The average standard deviation of the height is 4.6 μm for a single star.

$s_{\rm rms}$, is the sum in quadrature of $s_{\rm rms,min}$ and the linear dependence of the spot size on the defocus value $(z - z_0)$. For relatively small defocus values, a taylor expansion shows that $s_{\rm rms}$ follows an approximately parabolic shape. In fact, we define the parabolic focus range as that over which this approximation is valid,

$$s_{\rm rms} = \langle \sqrt{{\rm K}(z - z_0)^2 + s_{\rm rms,min}^2} \rangle \propto {\rm K}' + \frac{1}{2}(z - z_0)^2 \tag{2}$$

for arbitrary constants, K and K′, which incorporate the F/# of the beam, the observing wavelength, and wavefront static RMS error. The important form of this equation is that it is symmetric in defocus value, meaning that z_0 is, to first order, independent of the physics incorporated into K′. Thus, simply by plotting and fitting a parabolic curve to the relationship between spot radius (or FWHM) and z, z_0 (a.k.a., the height at the star's position) can be obtained.

3.1.2 Sources of bias

Because the parabola method relies on the numerical best fit determination of the (local) parabolic minimum, the sources of focus error, or bias, are only those that result in a shift of this minimum location. Effects on the measured image size that are symmetrical in the defocus value (defined as $z - z_0$ in Sec. 3.1.1) are irrelevant to the determination of the minimum location. Spherical aberration, for example, results in an increase in image width that is symmetric on both sides of focus (i.e., z_0), and therefore does not impact the accuracy of the parabola method. Similarly, vignetting alone introduces a bias into the measured image width as a function of field position due to the finite primary mirror diameter. However, this impact is symmetric with respect to the defocus value. Because the same vignetting bias applies on both sides of focus, and is constant for any given star position in the field, the measured parabola minimum remains unaffected. Note, the vignetting due to both the centrally obscuring cryostat and support shadows, which also depends on field location, has no first-order impact on the image size at all (in the absence of aberrations), as long as the central obscuration remains interior to the entrance pupil.

However, in the combined presence of both spherical aberration and vignetting, second-order biases impact the parabolic minimum location (i.e., z_0) in the parabola method. The details of the combined aberrations are complicated, and we attribute the deviations from the ideal parabolic shape in FWHM vs. z to the combined aberrations. The altered shape of ZTF's FWHM vs. z can be seen by comparing how well the best parabolic fit matches the observations for a single star in ZTF (Fig. 2) to PTF (Fig. 3), because the vignetting and spherical aberrations are significantly smaller in PTF. The altered shape of the ZTF observations results in increased uncertainty in the coefficient describing the parabolic minimum location (z_0) returned by the fitting algorithm (see Step 6 in Sec. 3.1.3). The standard deviation of z_0 for a given star increases on average from 4.6 to 15 μm with PTF and ZTF, respectively. Note that the deviations from the parabolic shape might be exacerbated by assuming a two-dimensional Gaussian PSF when measuring the FWHM of the stars.

Changes in atmospheric seeing during a through-focus scan could systematically increase or decrease the FWHM measured for a given star as z is changed, which would alter the z_0 measured. To adjust for this, we randomize the order we took the images, instead of steadily increasing or decreasing the z position. This way changes in seeing would increase the uncertainty when fitting for z_0, but would not bias it to a different z position.

3.1.3 Parabola procedure

Now that we understand the algorithm and biases, here we describe the detailed procedure used to measure the height for each section of the SFP using the parabola method.

Step 1: Collect images. We experimented with taking images of the same stars with the hexapod z position changing by a fine step size of 0.04 mm steps (as seen in Fig. 2), but found that a step of 0.1 mm was sufficient. We need images 0.2 mm on either side of the highest and lowest section of the SFP in order to span the necessary z range to adequately measure the height of every section. In May on ZTF, this was achieved with nine 30 sec images from -8.92 to -9.72 mm with 0.1 mm steps, where the optimal hexapod z position (i.e., the z that places the majority of the SFP in focus) was -9.4 mm.

Step 2: Clean images. Images were sufficiently clean without bias or flat field corrections, so we only needed to use PyRAF[†] to correct for pre-scan and over-scan regions.

Step 3: Identify stars and measure properties. We used SExtractor[3] to measure the x and y coordinates (in pixels) and the FWHM of the stars seen in each image. In order to apply S/N cuts we used SExtractor to measure the flux with a variety of different apertures and the background. We also had SExtractor return the elongation parameter to remove streaks.

Step 4: Purify star sample. To only include the cleanest sample of stars when measuring the height per SFP section, we reject the following groups of stars. We exclude all stars within 25 pixels of a CCD edge to avoid edge effects. We reject stars with elongation values greater than 3.0 to ensure all streaks are removed. We calculate the S/N for each star for a range of apertures to account for different size stars, and select the aperture with highest S/N. We required the S/N to be greater than 50 to reduce the computation time of the next step. We also required the S/N to be less than 1,000 to avoid saturated stars. Finally, we rejected stars with outlier FWHM values (partially to avoid double stars), which were identified through four iterations of rejecting stars that have FWHM that are $3\sigma_{\rm MAD}$ outliers.

Step 5: Match stars in all images. We find stars that appear in all images and record the FWHM in each image. To reduce the computation time of this step, we break up the stars into the 2,304 SFP sections and only search for matches within a given section. Stars are considered a match if they are within 12 pixels of eachother in different images, and stars are rejected if more than one match is identified in an image.

Step 6: Find height for each star. We fit a parabola to the FWHM vs. z observations for each star to find z_0 (a.k.a., position of the parabolic minima or the height for that star) and its error—see Fig. 2. The parabolic fit is done using Python's scipy.optimize.curve_fit[4] function, where $\mathrm{FWHM} = \beta(z - z_0) + \mathrm{FWHM}_0$ was fit to the observations. This fit returned the optimal coefficients (β, z_0, and FWHM_0) and the corresponding estimated covariance matrix. The square root of the diagonal values of the covariance matrix gives the standard deviation of each coefficient. Therefore, this fit returns z_0 and the standard deviation of z_0 for a given star.

[†]PyRAF is a product of the Space Telescope Science Institute, which is operated by AURA for NASA.

Once z_0 is measured using all the images in the through-focus scan, we re-measure z_0 with a subset of the images. We do this to avoid becoming biased by regions far from focus where the parabolic approximation is no longer valid (see Eq. 2). The subset of images is defined as the closest five images to the initial z_0 measured or all images within 0.2 mm of z_0, whichever includes more images. Using this subset we measure the final z_0, which gives the height for each star.

Step 7: Average height for each SFP section and measure precision. The heights of all of the stars with x and y positions within a given SFP section are averaged together to determine the height for the section. The heights of the stars are averaged using a weighted average that weights the stars by the reciprocal of the height variance, where this is the error returned from the parabolic fit mentioned in the previous step. If we average the standard deviation of the height for all ZTF stars in the SFP, we measure a standard deviation of 15 μm, which is the typical height error per star. If we average the standard deviation of the height for all ZTF sections in the SFP, we measure a standard deviation of 9 μm, which is the typical height error per SFP section. We conclude the parabola method with ZTF measures the SFP height with a precision of 9 μm.

Step 8: Measure the stability of the height map. We repeat this entire analysis five times with observations from three different days. After removing the global tip and tilt from each height map (see Sec. 4.1), we average all five maps to create a master height map determined using the parabola method (look ahead to Fig. 13 to see this master map). To quantify the stability of the height map, we subtract each of the five height maps from the average height map and see how the measured heights of each section changes. For each of these five 'difference maps' we measure the difference between the section that had the largest increase in height and the section that had the largest decrease in height and report this as the peak-to-valley difference. The peak-to-valley difference between each of the five height maps and the averaged map is [40.6, 67.2, 37.4, 25.5, 33.8] with an average of 41 μm. The standard deviation of the height changes between an individual map and the average map is [6.3, 10.2, 6.5, 4.6, 5.8] with an average of 6.7 μm. The peak-to-valley difference shows that the height in a given section can have large fluctuations. However, the average standard deviation of the height difference between maps of 6.7 μm is consistent with the height uncertainty of 9 μm per section reported in Step 8. We conclude that the ZTF height map is stable within our measurement height errors of 9 μm per SFP section.

3.1.4 Test with PTF to ensure precision

In October 2015, we tested the parabola method on the PTF system as both a verification that this method can achieve the required precision and check against an existing map measured using a profilometer in 2008.[5]

We used the same procedure described in Sec. 3.1.3 with some small differences. In addition to pre-scan and over-scan corrections, PTF images need bias and flat field corrections. Each 4k×2k CCD in PTF is divided into $\sim 410 \times 270$ pixel sub-frames for a total of 75 sections per CCD. Two CCDs are missing from this analysis—one was temporarily not functioning properly (the top left CCD in Fig. 4) and one wasn't functioning for the duration of the survey.

Smaller spherical aberrations and vignetting with PTF resulted in smaller height uncertainties. The average standard deviation of the height for a star was 4.6 μm (see Fig. 3). You can see a height map for individual stars in Fig. 4. The map of the height for the CCD sections is shown in Fig. 5, where the average standard deviation of the height for a CCD section was 1 μm.

To measure the stability of the height map, we used three different maps measured from three through-focus scans from the same night. Subtracting these three maps showed that the height differences for CCD sections ranged from −15 to +15 μm, so a peak-to-valley differences of 30 μm. The height measurement error of a CCD section (1 μm) is well below the height differences seen in different maps ($\lesssim$ 30 μm). It is possible that we are seeing the PTF CCD mosaic changing shape within a night, possibly due to changing environment (e.g., thermal fluctuations, gravity, or seeing dependent biases). However, additional analysis would have been needed to confirm this. Although the maximum height change between different height maps is larger than the ZTF requirement of 10 μm, more than 90% of the CCD sections had height differences less than 10 μm. Given the measurement error for the height of a CCD section was 1 μm and more than 90% of the CCD sections were stable within 10 μm, this test confirmed that the parabola method would be able to meet ZTF's requirement.

As an important sanity check we compared the height map we measured in October 2015 using the parabola method with the height map published in 2008 using profilometry[5] (compare Figs. 5 and 6). The maps measure

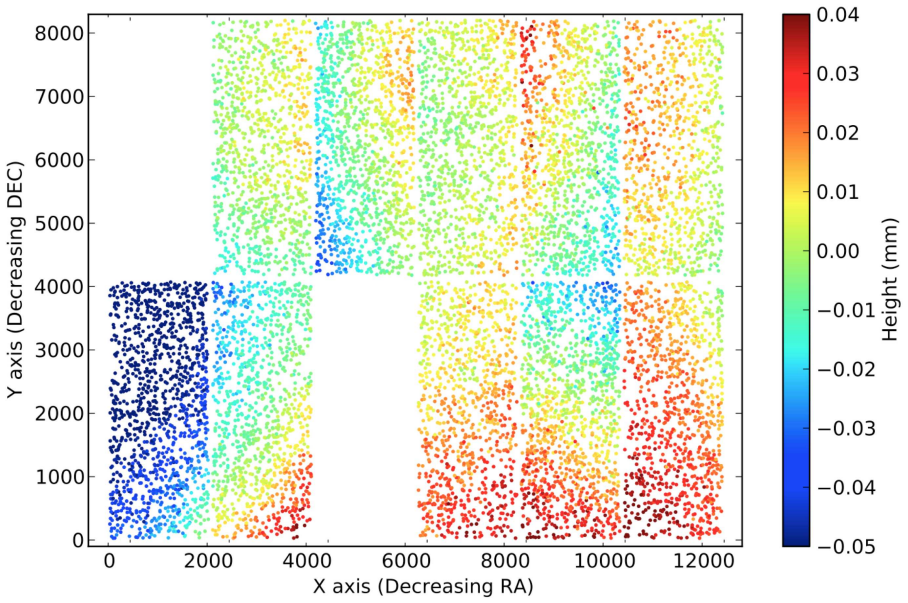

Figure 4. Height map of PTF's SFP created from images taken in October 2015. This map shows the heights of individual stars, which have an average height uncertainty of 4.6 μm standard deviation.

heights of the same overall spatial pattern and the same range across PTF's SFP of $\sim 90\,\mu$m. Height differences between the two maps of up to $\sim 30\,\mu$m are seen when comparing specific CCD sections, which is same height difference that was observed when comparing the three PTF height maps to each other. It is worth noting that differences could be due to curvature introduced by the optics which the parabola method would detect but the profilometer would not, simply because the profilometry was measured in the lab instead of in situ. We conclude that the height maps created from profilometry and the parabola method are sufficiently in agreement, and consider this a successful sanity check.

3.2 Donut method: second moments of "donuts"

An alternative method of determining the height of the SFP relies upon the calculation of the second moment of image intensity from images sufficiently out of focus such that the stars appear as blurred donuts. We can theoretically use this donut method with a single defocused image; however, it is subject to errors such as those induced by atmospheric seeing. Therefore we use images on either side of focus to measure the height of a SFP section.

3.2.1 Donut algorithm

Given defocused images on both sides of focus, at hexapod z positions F^- and F^+, one can derive the height at which a SFP section is in focus (F_0[‡]) in terms of the calculated second moments. Below, we summarize this method originally published in Ref. 6 and refer you to Fig. 7 for a graphic depiction.

For a given defocused and pixelated image intensity, I_{ij}, the second spatial moment in x and y are M_x and M_y, respectively, and are defined as

$$M_x = I_0^{-1} \sum (x_{ij} - x_c)^2 I_{ij} \text{ and } M_y = I_0^{-1} \sum (y_{ij} - y_c)^2 I_{ij}. \tag{3}$$

The zeroth and first moments used in Eq. (3) are

$$I_0 = \sum I_{ij},\ x_c = I_0^{-1} \sum x_{ij} I_{ij},\ \text{and } y_c = I_0^{-1} \sum y_{ij} I_{ij}. \tag{4}$$

[‡]Note that F_0 is identical to z_0 used in Sec. 3.1. F_0 is used here to be consistent with the equations in Ref. 6.

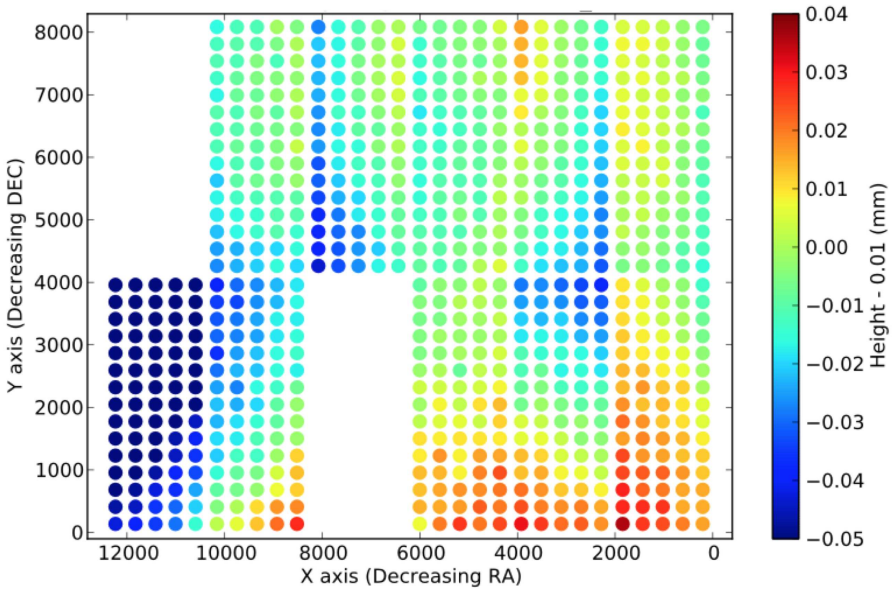

Figure 5. Height map of PTF's SFP created from images taken in October 2015. Each CCD is broken into 75 410×270 pixel sub-frames. This map shows the heights of sections, measured by averaging all stars that are in a given section. Each section has an average uncertainty of $1\,\mu$m standard deviation. We compared three maps with images taken during the same night and found height differences within the ZTF requirement ($10\,\mu$m) for 90% of the SFP sections. This verified the stability of the maps and validated the parabola method could be used for ZTF.

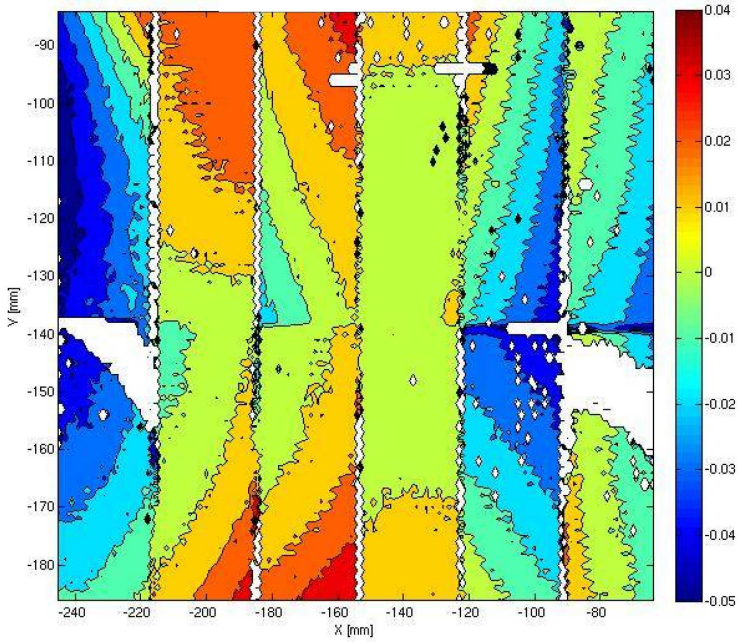

Figure 6. *Reproduced from Ref. 5, Fig. 9.* Profilometry height map published in 2008 in Ref. 5. After they subtracted the global tip and tilt of PTF's SFP, they found the map varied $90\,\mu$m peak-to-valley and $17\,\mu$m RMS. The calibration they performed reported an accuracy of $10\,\mu$m peak-to-valley and $1\,\mu$m RMS. Comparing this map to Fig. 5 provided an important sanity check, since the overall height pattern was consistent and the same peak-to-valley difference was observed across the plane.

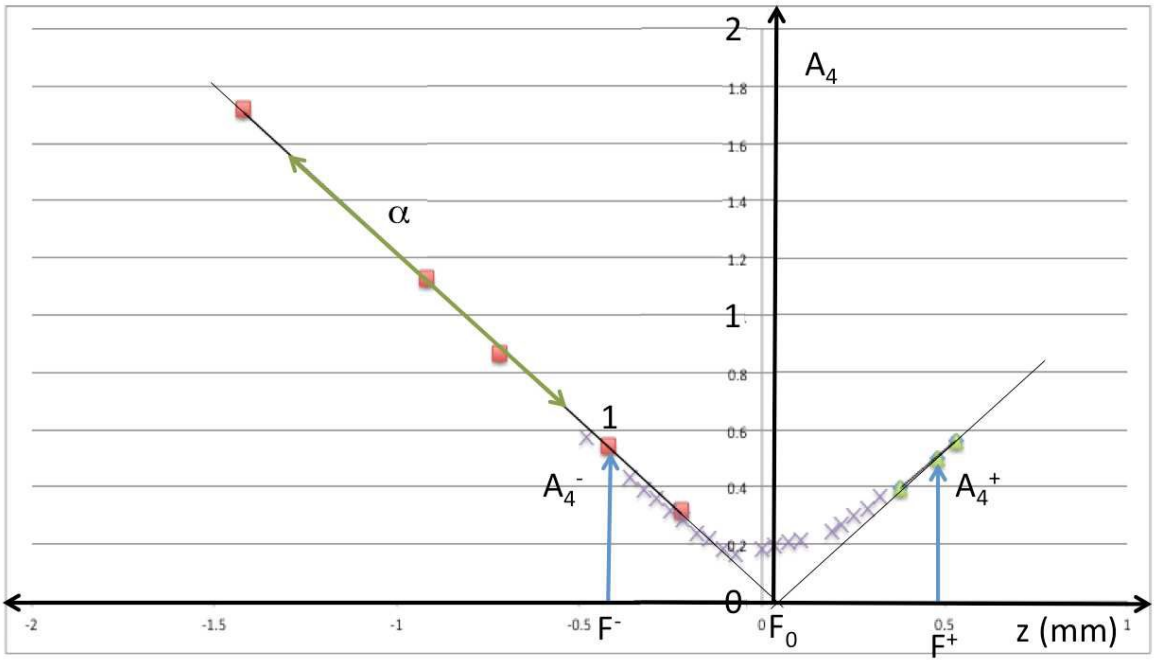

Figure 7. Demonstration of the donut method. This is how we use the donut method to determine the height at which a SFP section is in focus (F_0), which is described in Sec. 3.2.1. The A_4 parameter defined in Eq. (5) is plotted vs. the hexapod z position. The squares each represent a 120 sec image—red for images taken below focus and green for images above. The 'x's mark 30 sec images observed for the parabola method. Two thin black lines represent the lines in Eq. (6) with slopes of α that pass through an image below focus (A_4^- at position F^-) and above focus (A_4^+ at position F^+). They intersect at position F_0 with a value of $A_4 = \delta$. The δ parameter is essentially zero in this plot, so is not labeled.

We consider the second moments (M_x and M_y) together to define the parameter A_4,

$$A_4 = \mathrm{p}\sqrt{(M_x + M_y)/2}, \tag{5}$$

where p is the angular size of the detector pixels in arcseconds (1.0 arcseconds/pixel for the ZTF camera). With an image above focus (F^+, A_4^+) and below (F^+, A_4^+), two linear relationships are given:

$$A_4^- = \alpha(F_0 - F^-) + \delta \text{ and } A_4^+ = \alpha(F^+ - F_0) + \delta, \tag{6}$$

where α is the slope, F_0 is the z position where the lines intersect, and δ is the A_4 value at F_0 (see Fig. 7). Here we assume that α and δ are identical on both sides of focus; the effect of these assumptions are discussed in Sec. 3.2.2. The two A_4 vs. z lines can be combined to solve for F_0, yielding

$$F_0 = (F^+ + F^-)/2 + (A_4^- - A_4^+)/(2\alpha). \tag{7}$$

Note that determining the focus with the donut algorithm depends on knowing α. A global α parameter can be calculated for an instrument (e.g., ZTF) by using two images on the same side of focus; however, we find that a position-dependent α parameter is required to reduce error (see Secs. 3.2.2 and 3.2.3).

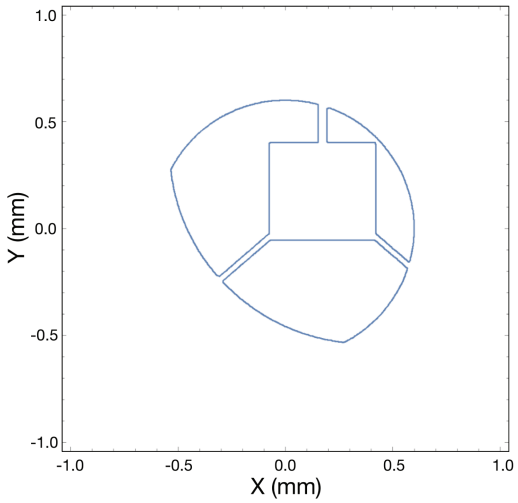

Figure 8. Example of the vignetted ZTF pupil used for image simulation. This pupil corresponds to an off-axis field point in the far corner of the ZTF CCD mosaic array (3.5, 3.5) degrees off-axis. The units of the abscissa and coordinate are in meters and are projected onto the telescope entrance pupil. We use the following values: the telescope pupil radius = 0.6 m, cryostat shadow width = 0.498×0.457 m, and primary mirror radius = 0.914 m.

3.2.2 Sources of bias

The donut method relies on the intersection of the linear trend of A_4 vs. z above and below focus, so—like the parabola method—the donut method is only impacted by biases that are asymmetric. Spherical aberrations and vignetting have symmetric impacts when treated individually (see Sec. 3.1.2 for a more in depth discussion of this).

When treated together, spherical aberrations and vignetting impact the determination of the height using the donut method, specifically by causing a significant dependence of α (see Eq. 7) on the star's position in the SFP. To confirm that the combined effect of spherical aberrations and vignetting is the origin of the dependence of α on position, we generate simulated image data[§] of the ZTF optical system including both $W_{040} = -7$ waves and realistic vignetting across the field (Fig. 8). To explore the accuracy of the donut method[¶], we generated synthetic image data for each CCD, under the model that individual CCDs sat upon a staircase of relative z positions, each positioned 25 μm higher than the last, in CCD numbering order with CCD S01 being the 'lowest' and thus requiring the largest $+z$ value of hexapod position to focus this CCD. The results of this analysis confirm that the position dependence of α is due to spherical aberrations and vignetting effects, and is particularly sensitive to the large vignetting present in ZTF. This motivated measuring an α map of the SFP from which a polynomial was fitted to determine α at a given pixel (see Step 5 in Sec. 3.2.3).

An additional impact of the combination of spherical aberrations and vignetting is that α differs slightly on either side of focus. As we mentioned when we introduced the two lines in Eq. (6), we assume that α and δ are identical on both sides of focus. A careful study of the two lines in Fig. 7 shows that this is not actually true. We believe α is not symmetric for the same reason the FWHM vs. z was not perfectly parabolic in the parabola method. The fact that α differs on either side of focus causes small errors when determining the z position where these lines intersect (i.e., F_0 or the height). Also note that assuming δ is the same on both sides is valid if the atmospheric seeing is constant for all images used.

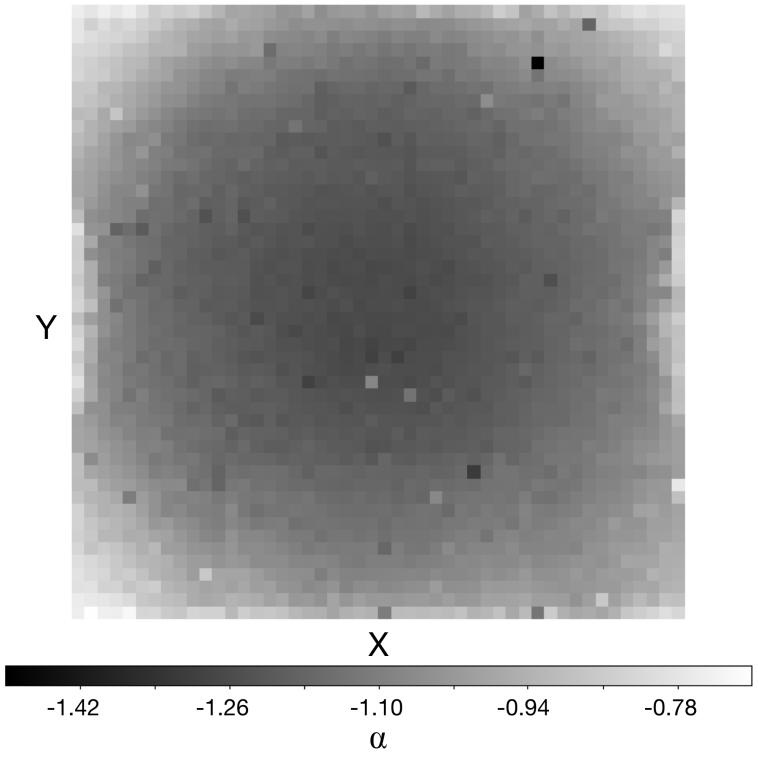

Figure 9. Master map of α as a function of position. This is a map of the α parameter as a function of pixel location for the entire SFP. The master map is the median combination of many α maps made with pairs of images on the same side of focus. A parabolic function is fitted to the map shown here, and this function is used when creating height maps with the donut method.

3.2.3 Donut procedure

Keeping these biases in mind, here is the detailed procedure we have designed to measure the height for each section of the SFP using the donut method.

Step 1: Collect images. We find that the optimal regime for A_4 values fall within 0.5 and 1.5 for the donut method. In this regime, sources have high S/N, sources are compatible with SExtractor alogirthms, and we are in the linear regime of A_4 vs. z. For ZTF, this corresponds to images roughly 0.5 mm out of focus. We obtained a large set of images taken at the following hexapod z positions: -8.0, -8.5, -9.0, -9.2, -9.8, -9.9, and -9.95 mm. Note that for this set of images the exposure time was 120 seconds and the optimal hexapod z position was -9.41 mm. However, we find that four images—two on each side of focus—roughly 0.5 mm out of focus are needed to calculate a height map. Two images on the same side of focus are used to measure a map of α and two images on different sides of focus are used to measure a height map of the SFP.

Step 2: Clean images. All images are processed using IRAF[7] by removing the pre-scan and over-scan regions, subtracting a master bias image, and dividing by a master flat field image.

Step 3: Identify stars and measure properties. We extracted the sources from each of the images using SExtractor to obtain the x and y coordinates (in pixels) and the second moments in x and y for each source. Default SExtractor parameters are not designed to find out of focus donuts, so a few adjustments are necessary. The following SExtractor parameters should be adjusted: DETECT_MINAREA, DEBLEND_MINCOUNT,

[§]The code to generate this vignetting can be found in the Mathematica workbook, "Analytic moment for obscured image.nb".

[¶]Due to the computational burden of producing these simulations, producing enough images to constrain the parabola method was deemed unnecessary.

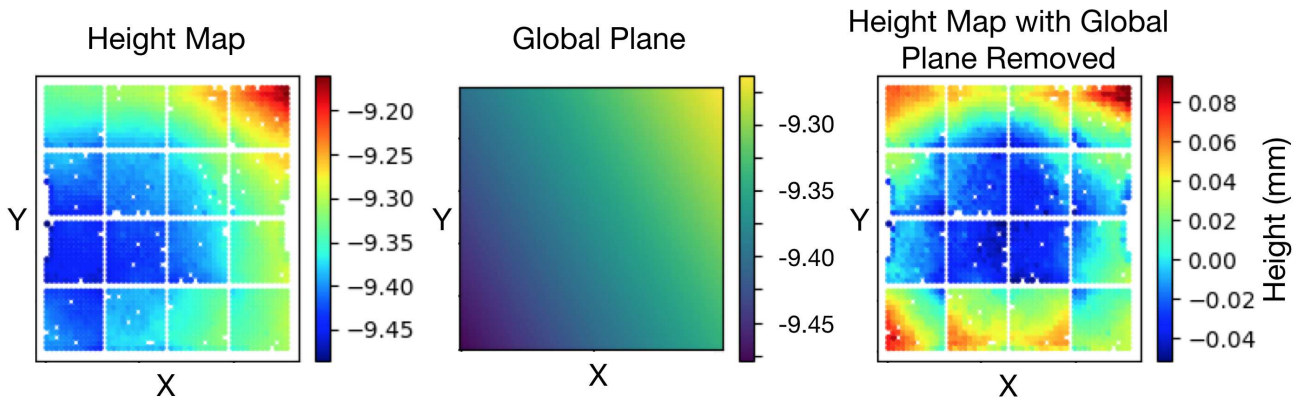

Figure 10. Removing the global tip/tilt from a height map. A global plane (center) is fitted to the original height map (left). The plane is subtracted from the original height map to create a new height map (right) that has no global tip/tilt remaining. This new height map can now be accurately combined or compared with other height maps.

DETECT_THRESH, and ANALYSIS_THRESH[‖]. When processing ZTF images the object and memory stack would often over-flow causing some sources to fail during extraction. To adjust for this we increased the MEMORY_OBJSTACK, MEMORY_PIXSTACK and MEMORY_BUFSIZE by a factor of three from the defaults.

Step 4: Measure median value of $\mathbf{A_4}$ for each SFP section. The original implementation of the donut algorithm (Sec. 3.2.1) was very computationally intensive. Attempts to identify the same star in images on both sides of focus yielded many stars that could not be properly matched and decreased the population of analyzed stars dramatically. This was a failure of the cross-correlation algorithm and not a fundamental problem with the process. Instead of matching individual donuts, we determined the median value of the A_4 metric for each section of the SFP.

Step 5: Measure $\boldsymbol{\alpha}$ as a function of position. In order to calculate the focus from Eq. 7 we need to determine α. We can measure α as a function of position by determining the slope of the line when plotting A_4 vs. z for two images on the same side of focus. Each pair of A_4 images on the same side of focus produces a single α map, and a set of α maps can then be combined into a single master α map (see Fig. 9). The master α map can then be fitted to a parabolic function that is then used to calculate α as a function of x and y pixel coordinates. The following parameters describe this function: α =p00+p10x+p01y+p20x^2+p11xy+p02y^2[**].

Step 6: Find height for each SFP section. Once α has been determined, any pair of images taken on opposite sides of focus can be used to calculate the height from the measured A_4^+ and A_4^- donut metric (based on the SExtractor determined second moments). Height maps were constructed using all of the possible combinations of images taken on opposite sides of focus—11 pairs were used in total. After selecting a set of high quality height maps and removing the global tip/tilt from each map (see Sec. 4.1), the height maps were combined using IRAF (combine=median) to create a master height map (skip ahead to Fig. 13 to view this master height map).

4. MEASURING ADJUSTMENTS TO THE SCIENCE FOCUS PLANE FROM HEIGHT MAPS

4.1 Remove global tip/tilt from each map

Althought we did flatten the SFP by iteratively shifting the hexapod and creating a FWHM map (see Fig. 1 and Sec. 2), global tip/tilts need to be measured and removed from each height map before comparing or averaging them. These global tip/tilts occur due to incorrect alignment of the hexapod and other effects—such as shifting caused by the telescope moving to or being at different positions. Different height maps measured at the same

[‖]For ZTF images we used the following values: DETECT_MINAREA=100-600, DEBLEND_MINCOUNT=0.5, DETECT_THRESH=3.0, ANALYSIS_THRESH=3.0.

[**]For ZTF this equation is $\alpha = -0.7161 - 4.387\times10^{-5}x - 4.496\times10^{-5}y + 1.693\times10^{-9}x^2 + 1.114\times10^{-10}xy + 1.719\times10^{-9}y^2$.

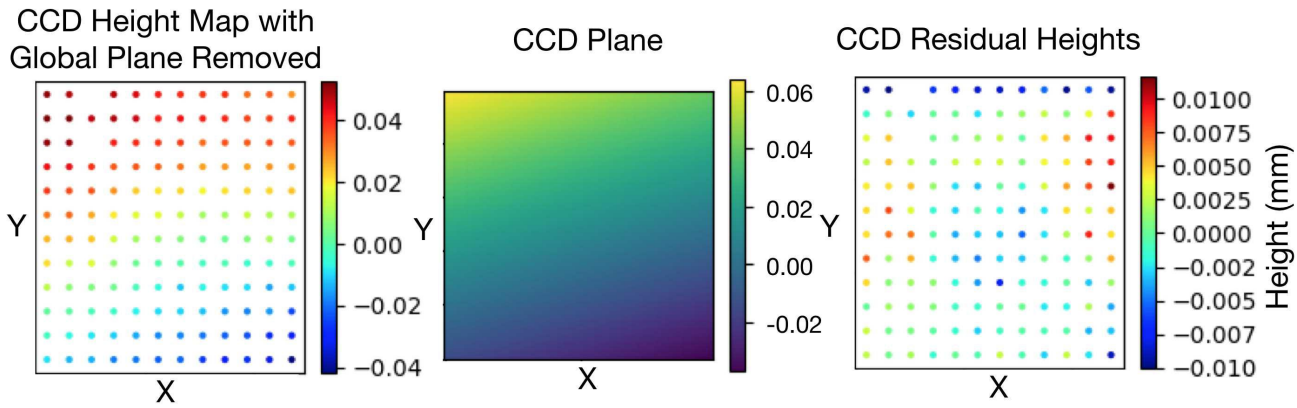

Figure 11. Fit a plane to each CCD. After the global tip/tilt has been removed from a height map, we can fit a plane (center) to each CCD height map (left, CCD S02 in this example). This plane provides a formula to calculate the height corrections that should be applied to each CCD support in order to flatten the SFP (see Table 1). Subtracting the plane from the CCD height map results in the CCD residual heights (right).

reported hexapod position show different global tip/tilts, so incorrect alignment of the hexapod clearly is not the the only source of misalignment. The global tip/tilt is measured by fitting a plane, $z = ax+by+c$, to the height map of the entire SFP. Subtracting this plane from the height map removes any global tip/tilt.

The height maps measured in Sec. 3 are formatted as a list of CCD #, x, y, and z_0 values for each SFP section. The x and y parameters are the pixel coordinates for the center of the CCD section with the bottom left pixel of each CCD having coordinate (0,0). The z_0 parameter is height of the section (i.e., the hexapod z position where this section of the SFP is in focus). To fit a plane to the entire SFP, the x and y section coordinates need to be transformed to place them relative to the entire SFP instead of relative to each CCD. We also properly account for the chip gap between CCDs by transforming the distance into number of 15 μm pixels that would fit in the chip gap. For ZTF, there is a 462 pixel-sized gap between CCDs in the x direction and a 645 pixel gap in the y direction. Once we have x and y positions that place the CCDs in the correct pixel positions (see left panel in Fig. 10), we fit a plane to the entire SFP.

We fit a plane to a height map using a least squares fit to the data (see center panel in Fig. 10). This is done either by using Python's astropy.modeling[8††] for height maps created by the parabola method or Matlab's Curve Fitting Toolbox for height maps created by the donut method.

At this point we subtract the plane from the data to return the height map with the global tip/tilt removed (see right panel in Fig. 10). Switching the x and y coordinates back to the values relative to each CCD prepares the height map for the next stage—fitting a plane to each CCD.

4.2 Fit a plane to each CCD

Once the global tip/tilt is removed from the SFP height map, the height maps from each method can be combined (see Step 8 in Sec. 3.1.3 and Step 6 in Sec. 3.2.3). When we have a master height map for each method, we fit a plane ($z = ax+by+c$) using the same technique described in Sec. 4.1 to the height map of each CCD (see left panel in Fig. 11). The coefficients of the plane equation (a, b, c) are then used to calculate a new image that represents the plane fit to the individual CCD (see center panel in Fig. 11). The resulting plane fit image is then subtracted from the master height map to evaluate the residual features in the image that are not accounted for by the planes (see right panel in Fig. 11).

As an aside, this technique can be used to study the shape of the individual CCDs. Figure 12 shows the residual features of each CCD after the CCD plane is subtracted. This shows $\sim 8-15\,\mu$m peak-to-valley height

††We fit a plane using Python with these two commands: p_init = astropy.modeling.models.Polynomial2D(degree=1) and fit_p = astropy.modeling.fitting.LevMarLSQFitter()

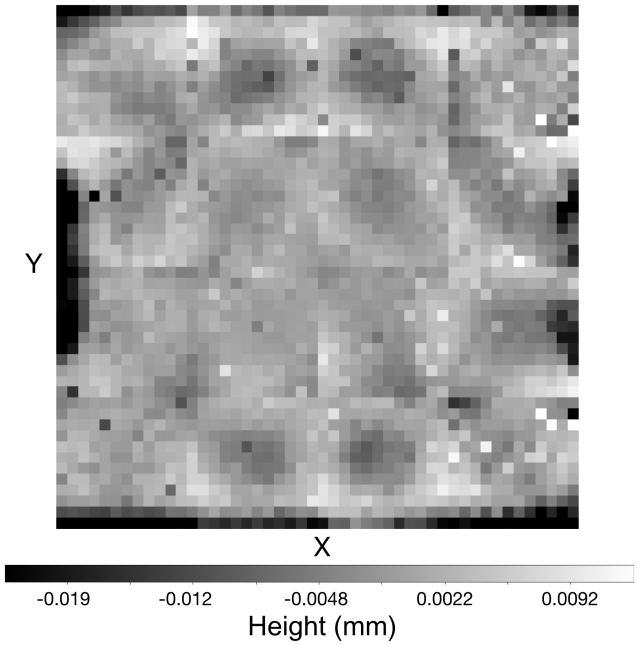

Figure 12. CCD residual heights. This is a map of each CCD's residual heights (an example can be seen in the right panel of Fig. 11). Although CCD gaps are not displayed in this image, the 16 CCDs can still be distinguished by eye. This height map shows that many of the CCDs have a concave shape and have $\sim 8-15\,\mu$m height differences. Please note that this image includes contamination due to outgassing material (specifically see the darkened edges). The gaskets that support the cryostat window have since been replaced and all signs of contamination have been removed.

differences within a CCD with a typically concave shape. It is important to note that these residuals also include light contamination (e.g., scattered light), so this is not a pure measurement of CCD shape. For these ZTF images, you can see contamination from outgassing material. Since these images were taken, all outgassing material has been replaced and all signs of contamination have been removed.

5. COMPARE THE RESULTS OF THE PARABOLA AND DONUT METHODS

The master height maps measured using the two different methods—parabola and donut—are shown in Fig. 13. The two methods produced very similar height maps—both in magnitude and the overall spatial pattern. By subtracting these two height maps (Fig. 14), we measure a maximum section height decrease of $-22\,\mu$m and a maximum section height increase of $+26\,\mu$m, which results in a $48\,\mu$m peak-to-valley height difference between the maps. The standard deviation of the height differences between these two maps is $8\,\mu$m. If you recall, the precision that we were able to calculate the height of a section of the SFP using the parabola method was $9\,\mu$m (Step 7 in Sec. 3.1.3), so this level of difference between the two maps is within our measurement error. In addition, similar height differences were observed between maps measured solely by the parabola method ($41\,\mu$m peak-to-valley and $6.7\,\mu$m standard deviation, Step 8 in Sec. 3.1.3). We do not see a clear reason why one method is more intrinsically precise than the other, so our final results are an average of these two maps. From this average, we now can determine how the SFP should be repositioned.

6. FINAL SCIENCE FOCAL PLANE ADJUSTMENTS

We fit the final CCD planes to each CCD from the final height map—the average of the master height map created via the parabola and donut method. The coefficients for each CCD plane is given in Table 1. The CCD naming convention is that the CCD number increases as you follow the row from left to right, and then continue this trend on the row below. Therefore the top-left CCD is S01, the top-right CCD is S04, the bottom-left CCD is S13, and bottom-right CCD is S16.

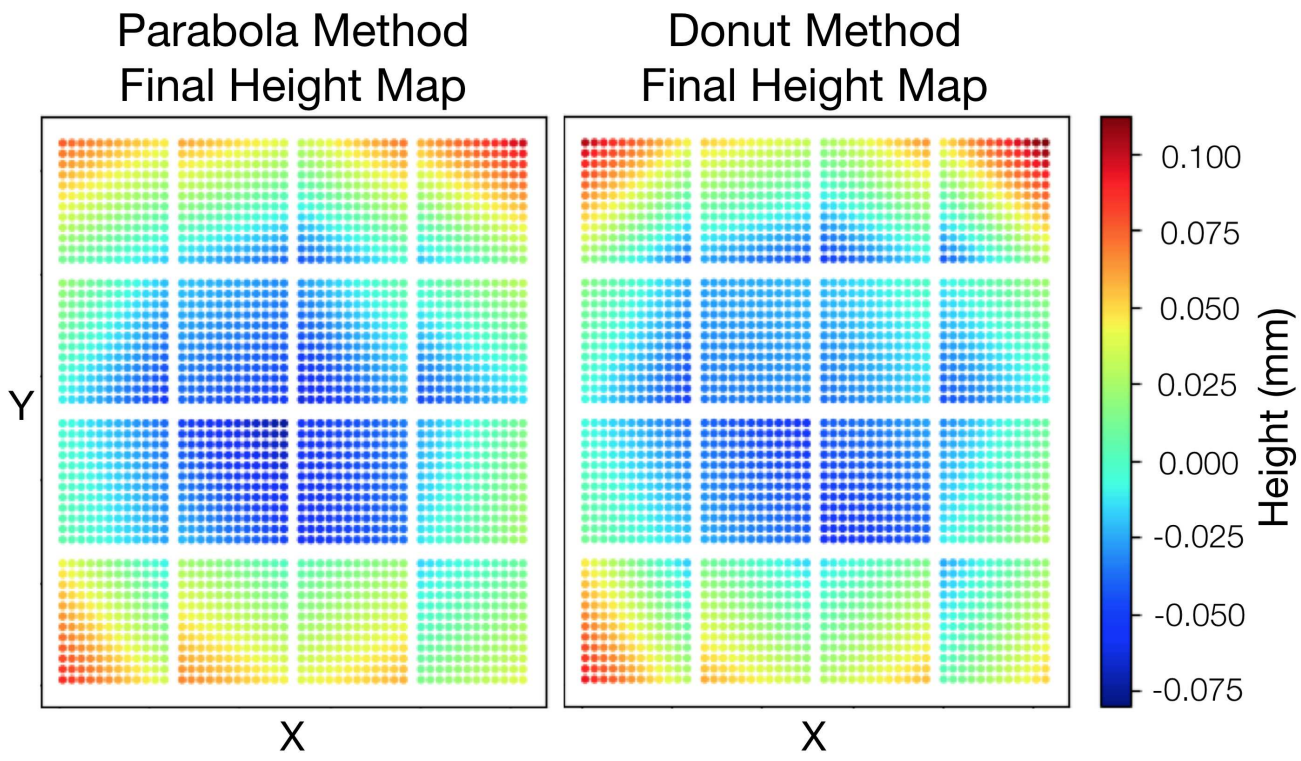

Figure 13. Master height maps using the parabola method (left) and donut method (right). This shows the CCD plane fitted to each CCD (see center panel in Fig. 11) for the two different methods. Figure 14 shows the difference between these two maps. An average of these two maps was used to calculate the final plane coefficients in Table 1.

The sign convention is that these coefficients represent the measured *hexapod z* position that brought each CCD into focus (with positive values moving the hexapod toward the telescope's primary mirror). Thus, these values also represent *the required adjustment* to the current CCD positions. For example, a negative fit plane height represents a point on the CCD where the hexapod had to move away from the primary mirror to reach focus, so this point on the CCD is currently too 'high' relative to the cryostat body and needs to be 'lowered' with respect to the cryostat body.

These planes therefore describe the correction that should be applied to the CCD supports in order to flatten the SFP. Each CCD has three supports and the height adjustment applied to each support can be found by finding the value of the CCD plane at the pixel above each support. The details of this conversion are beyond the scope of this proceeding.

7. CONCLUSION

The ZTF requirement is to measure the SFP within 10 μm RMS. Initially we use a stellar FWHM map to place the SFP in approximate focus. In variable seeing this technique can determine the hexapod z position where the SFP is in focus within 35 μm (Fig. 1). Following this, we use two different methods for creating a height map of the SFP—the parabola and donut method. Using the parabola method, the precision of the section heights is 9 μm standard deviation (Step 7 in Sec. 3.1.3), and the stability of the section heights is 6.7 μm standard deviation (Step 8 in Sec. 3.1.3). After removing the global tip/tilt from the height maps, we combine them to make a master map for each method. At that point we fit a plane to the CCDs in each map. The section heights of these two sets of planes (Fig. 13) agree within 8 μm standard deviation (Fig. 14), which is within the ZTF requirement. Our final height map is the the average of the results from the parabola and donut methods. The final CCD plane coefficients (Table 1) determine how the CCDs should be repositioned. Based on the

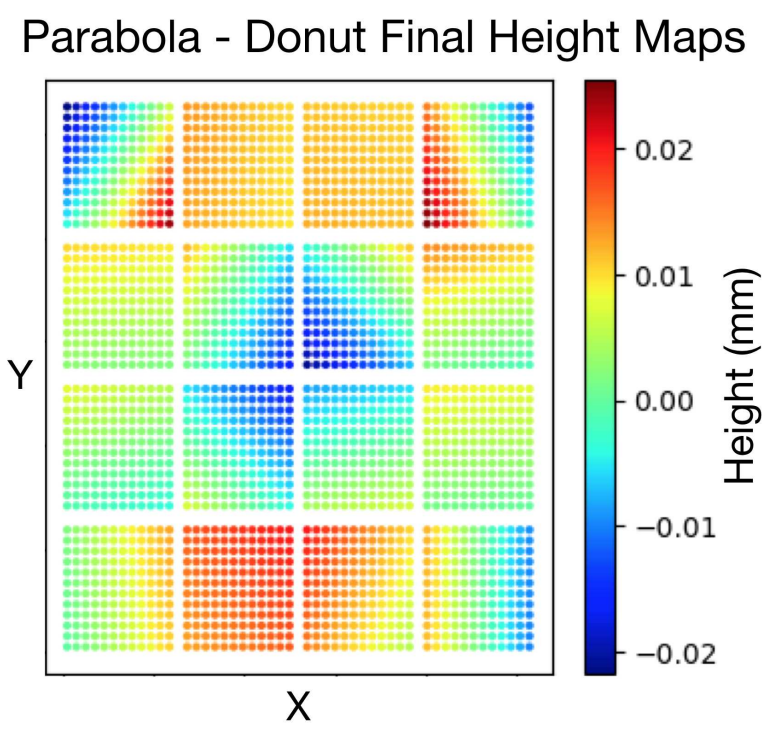

Figure 14. Height difference between the two methods. This is a map of the subtraction between the master height maps calculated using the parabola and donut methods (these maps are shown in Fig. 13). The height differences vary 48 μm peak-to-valley and the standard deviation of this difference map plotted is 8 μm. The height difference of the results of our two methods (8 μm standard deviation) is within the ZTF requirement.

Table 1. Final plane coefficients for the 16 ZTF science CCDs. All z values describe the required adjustment to the current CCD positions to flatten the SFP. The equation of each plane is $z = ax+by+c$, where a positive height value implies that point on the focal plane needs to be raised (moved in $+z$ with respect to the cryostat cold plate). The origin for the (x,y) coordinate system is the center of lower-left pixel (0,0) for each CCD.

CCD	a (μm/pixel)	b (μm/pixel)	c (μm)
S01	-0.00772	0.01374	16.9
S02	-0.00289	0.01409	-18.4
S03	0.00714	0.01368	-40.9
S04	0.01018	0.01581	-29.8
S05	-0.00728	0.00538	-9.6
S06	0.00026	0.00315	-39.8
S07	0.00530	0.00509	-49.3
S08	0.00570	0.00711	-40.2
S09	-0.00519	-0.00204	2.1
S10	-0.00285	-0.00269	-27.5
S11	0.00343	0.00227	-58.2
S12	0.00731	-0.00215	-12.4
S13	-0.00921	-0.00772	81.8
S14	-0.00409	-0.00750	58.7
S15	0.00483	-0.00581	25.6
S16	0.00692	-0.00563	7.4

results presented in this paper, the CCD supports were reshimmed to adjust the CCD heights to flatten the SFP. Some errors were introduced when converting the planes described in Table 1 to height corrections to the CCD supports, and efforts to correct this are ongoing. Even with these flaws, the standard deviation of the height across the entire SFP measured before and after the SFP was repositioned decreased from 33 to 14 μm. We calculate that this reduction in the height variation would improve the typical FWHM from approximately $2.80''$ to $2.08''$. This decrease in FWHM would reduce the integration needed to reach a given brightness by 45%. Without a thorough calculation that incorporates all contributions to the survey speed (e.g., overhead time between images), we approximate the impact this has on the survey speed by converting the 45% time reduction directly to an increase in survey speed of 81%.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE1745301. This work is based on observations obtained with the Samuel Oschin Telescope 48-inch and the 60-inch Telescope at the Palomar Observatory as part of the Zwicky Transient Facility project, a scientific collaboration among the California Institute of Technology, the Oskar Klein Centre, the Weizmann Institute of Science, the University of Maryland, the University of Washington, Deutsches Elektronen-Synchrotron, the University of Wisconsin-Milwaukee, and the TANGO Program of the University System of Taiwan. Further support is provided by the U.S. National Science Foundation under Grant No. AST-1440341.

REFERENCES

[1] Dekany, R., Smith, R. M., Belicki, J., Delacroix, A., Duggan, G., Feeney, M., Hale, D., Kaye, S., Milburn, J., Murphy, P., Porter, M., Reiley, D. J., Riddle, R. L., Rodriguez, H., and Bellm, E. C., "The Zwicky Transient Facility Camera," in [*Ground-based and Airborne Instrumentation for Astronomy VI*], *Proceedings of the SPIE* **9908**, 99085M (Aug. 2016).

[2] Reiley, D. J., Dekany, R. G., Smith, R. M., Delacroix, A., Feeney, M., and Callahan, S., "Optical design of the Zwicky Transient Facility: a major upgrade to the 48" Schmidt Camera at Palomar Observatory," in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **10590**, 105901U (Nov. 2017).

[3] Bertin, E. and Arnouts, S., "SExtractor: Software for source extraction.," *Astronomy and Astrophysics, Supplement* **117**, 393–404 (June 1996).

[4] Jones, E., Oliphant, T., Peterson, P., et al., "SciPy: Open source scientific tools for Python," (2001–). [Online].

[5] Rahmer, G., Smith, R., Velur, V., Hale, D., Law, N., Bui, K., Petrie, H., and Dekany, R., "The 12K×8K CCD mosaic camera for the Palomar Transient Factory," in [*Ground-based and Airborne Instrumentation for Astronomy II*], *Proceedings of the SPIE* **7014**, 70144Y (July 2008).

[6] Tokovinin, A. and Heathcote, S., "Donut: Measuring Optical Aberrations from a Single Extrafocal Image," *Publications of the ASP* **118**, 1165–1175 (Aug. 2006).

[7] Tody, D., "IRAF in the Nineties," in [*Astronomical Data Analysis Software and Systems II*], Hanisch, R. J., Brissenden, R. J. V., and Barnes, J., eds., *Astronomical Society of the Pacific Conference Series* **52**, 173 (Jan. 1993).

[8] The Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., Günther, H. M., Lim, P. L., Crawford, S. M., Conseil, S., Shupe, D. L., Craig, M. W., Dencheva, N., Ginsburg, A., VanderPlas, J. T., Bradley, L. D., Pérez-Suárez, D., de Val-Borro, M., Aldcroft, T. L., Cruz, K. L., Robitaille, T. P., Tollerud, E. J., Ardelean, C., Babej, T., Bachetti, M., Bakanov, A. V., Bamford, S. P., Barentsen, G., Barmby, P., Baumbach, A., Berry, K. L., Biscani, F., Boquien, M., Bostroem, K. A., Bouma, L. G., Brammer, G. B., Bray, E. M., Breytenbach, H., Buddelmeijer, H., Burke, D. J., Calderone, G., Cano Rodríguez, J. L., Cara, M., Cardoso, J. V. M., Cheedella, S., Copin, Y., Crichton, D., DÁvella, D., Deil, C., Depagne, É., Dietrich, J. P., Donath, A., Droettboom, M., Earl, N., Erben, T., Fabbro, S., Ferreira, L. A., Finethy, T., Fox, R. T., Garrison, L. H., Gibbons, S. L. J., Goldstein, D. A., Gommers, R., Greco, J. P., Greenfield, P., Groener, A. M., Grollier, F., Hagen, A., Hirst, P., Homeier, D., Horton, A. J., Hosseinzadeh, G., Hu, L., Hunkeler, J. S., Ivezić, Ž., Jain, A., Jenness, T., Kanarek, G., Kendrew, S., Kern, N. S., Kerzendorf, W. E., Khvalko, A.,

King, J., Kirkby, D., Kulkarni, A. M., Kumar, A., Lee, A., Lenz, D., Littlefair, S. P., Ma, Z., Macleod, D. M., Mastropietro, M., McCully, C., Montagnac, S., Morris, B. M., Mueller, M., Mumford, S. J., Muna, D., Murphy, N. A., Nelson, S., Nguyen, G. H., Ninan, J. P., Nöthe, M., Ogaz, S., Oh, S., Parejko, J. K., Parley, N., Pascual, S., Patil, R., Patil, A. A., Plunkett, A. L., Prochaska, J. X., Rastogi, T., Reddy Janga, V., Sabater, J., Sakurikar, P., Seifert, M., Sherbert, L. E., Sherwood-Taylor, H., Shih, A. Y., Sick, J., Silbiger, M. T., Singanamalla, S., Singer, L. P., Sladen, P. H., Sooley, K. A., Sornarajah, S., Streicher, O., Teuben, P., Thomas, S. W., Tremblay, G. R., Turner, J. E. H., Terrón, V., van Kerkwijk, M. H., de la Vega, A., Watkins, L. L., Weaver, B. A., Whitmore, J. B., Woillez, J., and Zabalza, V., "The Astropy Project: Building an inclusive, open-science project and status of the v2.0 core package," *ArXiv e-prints* (Jan. 2018).