



Vera C. Rubin Observatory
Data Management

Verifying LSST Calibration Data Products

Robert Lupton, Andrés A. Plazas Malagón, and Christopher Waters

DMTN-101

Latest Revision: 2025-03-28

DRAFT



Abstract

A description of plans for verifying LSST's Calibration Data Products. This document covers our approach to verification element LVV-57, addressing requirement DMS-REQ-0130.

Draft

Change Record

Version	Date	Description	Owner name
1	YYYY-MM-DD	Unreleased.	Robert Lupton
1.1	2022-08-05	Added PTC section	Andrés A. Plazas Malagón
1.2	2022-09-27	Added gain test: 6.2	Andrés A. Plazas Malagón
1.2.0	2022-09-27	Replaced “master” with “combined” and “aggressor” with “source”.	Andrés A. Plazas Malagón
1.3	2023-02-21	Added subsection on linearizer test	Andrés A. Plazas Malagón

Contents

1	Text from LVV-57	1
2	Introduction	1
3	Defect Maps	2
4	Biases	3
5	Dark Current	4
6	Monitor Amplifier Gains	5
7	Non-linearity	5
	7.1 Linearity correction verification.	6
8	Saturation Levels	7
9	Crosstalk	8
10	WRITE ME Flat Fields	9
	10.1 Broadband and Monochromatic Flats	10
	10.2 Contamination of Filters and Other Optical Elements	10
	10.3 Flats for Background Estimation	10
	10.4 Flats for Photometry	11
	10.5 Outputs	12
11	WRITE ME Fringe Correction	13
12	WRITE ME Filter Transmission Curves	14
13	WRITE ME Pixel Size Effects	15
14	WRITE ME Brighter-Fatter	16
15	WRITE ME Ghosts and Ghouls	16

16 Photon Transfer Curve	16
A References	17

Draft

Verifying LSST Calibration Data Products

1 Text from LVV-57

(DMS-REQ-0130 [LSE-61])

Specification:

The DMS shall produce and archive Calibration Data Products that capture the signature of the telescope, camera and detector, including at least: Crosstalk correction matrix, Bias and Dark correction frames, a set of monochromatic dome flats spanning the wavelength range, a synthetic broad-band flat per filter, and an illumination correction frame per filter.

Description:

For every calibration mode, prove that the data can be processed. Can be done with simulated data and that from the auxiliary telescope. Will need to be redone with real LSST camera data.

2 Introduction

In some cases the data taken for one purpose may be reused for another, but I have made no attempt to capture this here.

The tests will be developed using TS8 data where possible (in some cases the camera team may not have taken appropriate data), and repeated with auxTel, lsstCam on the CCOB (when appropriate data is available), comCam on the 8.4m telescope (if that happens), and then lsstCam on the 8.4m telescope.

- When "flats" are called for, unless otherwise noted, they should be taken using a stable c. 600-700nm broad-band light source (monochromatic is acceptable, but will require longer integrations). The illumination should be reasonably uniform; this means that it would be preferable to take the data with no filter in the beam, or with the wavelength chosen to be near a filter's central wavelength.

- I have not specified whether these calibration products should be produced for corner rafts as well as science rafts.
- Depending on our final model for how the image is modified by the detector and REB, we may choose to use a slightly different set of ISR stages than are listed here.

The sections labelled *WRITE ME* have been copied from my other calibration documents, so need to be changed from an algorithmic description to a V&V plan.

3 Defect Maps

We expect the camera team to deliver a list of bad (i.e., ones that are unusable) and suspect (i.e., ones that should raise a flag during processing) pixels.

The bad pixels should be ignored in all the tests which I describe below.

Once all the calibration products described below are validated, we will check for completeness of the defect maps.

Take a set of dome flat fields at a set of logarithmically spaced intensity levels starting at 500 DN and increasing by a factor of 2 (i.e., 500, 1000, 2000, ...); at least 5 exposures must be taken at all levels. The exact levels are not important, and need not be accurately known. A sufficient number of exposures at each level must be taken to measure the mean flux with 0.5% errors (except that a maximum of 20 exposures per level is required; this is only sufficient to measure the 500-count flux to less than 1%). Also take 3 pocket-pumped flats at each level.

The individual frames will be run through full ISR processing (including cosmic ray removal using an $N(0,1)$ Gaussian PSF). The flat field used should be appropriate to correcting surface brightness not QE. Combine using a median or 5-sigma-clip to make a "combined" frame; the combination may be done using standard LSST image stacking code. The pocket-pumped data should not be included in the combined frames, but should be median-combined separately.

N.b. Defect interpolation should be disabled in this processing.

Tests:

1. Divide each resulting combined frame by its median.
2. For each combined frame, confirm that the median level is 1.0 to within statistical noise (after masking known defects)
3. For each combined frame, divide by the standard deviation image, and confirm that the 5-sigma clipped standard is 1.0 to within statistical noise (after masking known defects)
4. for each pixel, check if all combined frames are consistent to within statistical noise; if not, confirm that they are included in the defect map. Also identify pixels in the defect map that are not identified.
5. Search the pocket-pumped flats for dark pixels. Confirm that they are all included in the defect map. Identify pixels in the defect map that are not found

4 Biases

Bias frames will be produced by taking the median of N_b (nominally 20) bias exposures.

Note that the choice of N_b must be made to measure the bias frame sufficiently accurately to not increase the effective read noise by more than 2.5%.

The individual frames will be overscan-corrected and trimmed, and then median or 5-sigma clip combined using standard LSST image stacking code. The resulting combined bias will then be ingested.

Tests:

1. Process an independent bias frame through the ISR including overscan correction and bias subtraction
2. Confirm that the mean of the result is 0 to within statistical error
3. Confirm that the 5-sigma clipped standard deviation of each amplifier is within 5% of the nominal readnoise, as determined by a robust measure of the noise in the serial overscan

4. Run a CR rejection on the result and confirm that the unclipped standard deviation is consistent with the 5-sigma clipped value.

5 Dark Current

Dark frames will be produced by taking the median of N_d (nominally 100) dark exposures, each of length T_d (nominally 30s), giving a total exposure of c. 3000s.

Note that the choice of N_d and T_d must be made to measure the dark frame sufficiently accurately to not increase the effective read noise by more than 2.5% due to the dark subtraction.

The individual frames will be processed through overscan, bias, and dark ISR processing, and cosmic ray removal using an $N(0, 1)$ Gaussian PSF before being median or 5-sigma clip combined; the combination may be done using standard LSST image stacking code.

Tests:

1. Process an independent dark frame through the ISR including overscan correction, bias subtraction, dark subtraction.
2. Confirm that the mean of the result is 0 to within statistical error
3. Confirm that the 5-sigma clipped standard deviation of each amplifier is within 5% of the nominal readnoise, as determined by a robust measure of the noise in the serial overscan
4. Run a CR rejection on the result and confirm that the unclipped standard deviation is consistent with the 5-sigma clipped value.
5. Process the 150 "even" and "odd" visits separately, and subtract the two resulting dark calibration frames.
6. Confirm that the mean of the difference is 0 to within statistical error
7. Confirm that the unclipped standard deviation of the difference is consistent with the dark current and readnoise (as measured by a robust measure of the noise in the serial overscan of the individual frames), as corrected by the correct combination of N_b , N_d , and T_d .

Repeat all these tests using $N_d/5$, $5 \cdot T_d$ dark exposures and compare the results. If the results are consistent (allowing for the slightly increased readnoise), use the smaller set of longer exposures for future generation of dark frames.

6 Monitor Amplifier Gains

We are unable to measure the gain as accurately as with Fe55 measurements, but we can check the relative gain sensitivity between the 16 amplifiers in a CCD and make rough measurements of the absolute gain. Note that the variation of gain as a function of intensity level is covered under "non-linearity".

Take pairs of flats with c. 20000 counts, and run the full ISR (probably; this means that we are measuring things relative to the calibs).

Tests:

1. Form the function $(I_1 - I_2)/(I_1 + I_2)$ whose mean value is $1/\text{gain}$. Note that this is sensitive to brighter-fatter so bin (or use BF correction code).
2. Compare the gain estimated from a pair of flat fields to the gain value stored in each amplifier, to monitor any temporal changes. Check that the relative fractional difference is within a threshold (default: 5%).
3. Measure the gain values for each amplifier that minimises discontinuities at amp boundaries.

7 Non-linearity

We expect the camera team to deliver non-linearity curves for every amplifier, and also the level above which these non-linearity curves may be unreliable.

Take a set of flat fields at a set of logarithmically spaced intensity levels starting at 500 DN and increasing by a factor of 2 (i.e., 500, 1000, 2000, ...). The exact levels are not important, but they must be well known; both the flux level and shutter time must be well measured. If necessary, we can open the shutter in the dark, turn on the light source, turn it off, and then

close the shutter. N.b. It may be necessary to modify or augment this set of intensity levels to successfully carry out the specified tests.

An alternative dataset is taken by opening the shutter, and using the CBP with a mask with reasonably large holes to produce spots on the focal plane at the same set of intensity levels as above, with at least one spot on every amplifier. After every exposure, the parallels on the CCD must be run to transfer enough rows that the successive spots do not overlap. The sequence should be repeated with the spots moved to different points on the amplifier. If necessary, this test can be carried out using masks with fewer holes by repeating the scans with the CBP rotated about its pupil to move the spots.

The individual frames will be processed through overscan, bias, dark, and flatfield ISR processing.

Tests:

1. For each flat exposure, divide by the (known) product of intensity and exposure time. For each amplifier calculate the median intensity, and fit a suitable functional form to measure the non-linearity.
2. Analyse these non-linear curves to determine the point at which the data cannot be reliably linearised.
3. For each "spot" exposure, make a robust measurement (e.g. median) of the flux in the flat-topped profile of each spot; the exact details of what defines the top of the spot are TBD. Divide these value by the known fluxes delivered to the spots (as provided by the CBP's photodiode), and fit a suitable functional form to measure the non-linearity.
4. Analyse these non-linear curves to determine the point at which the data cannot be reliably linearised.

7.1 Linearity correction verification.

Test:

1. Given a test linearizer of a certain type (e.g., Spline, Quadratic, Polynomial, or Lookup Table) created with the Calibration Products Production pipelines, a Photon Transfer

Curve should be created with data that has undergone Instrument Signature Removal using this input linearizer. Then, a second linearizer should be produced using this PTC data. Linearity residuals of this second linearizer should be less than a predetermined threshold, depending on the linearizer type:

- (a) Spline linearizer: the maximum of the absolute value of the spline residuals should be less than a nominal threshold (default: 1%).
- (b) Quadratic linearizer: the quadratic non-linearity coefficient, c_0 , should be less than a nominal threshold (default: $1e^{-6}$).
- (c) Polynomial linearizer: the quadratic term should be less than a nominal threshold, as in the Quadratic linearizer case, and the higher order polynomial terms should be less than an array of nominal thresholds scaled by the quadratic term threshold.
- (d) Lookup Table: the maximum of the absolute value of the linearizer residuals (defined as the ratio of the corrections provided by the table and the flux level for each correction) should be less than a nominal threshold (default: 1%).

8 Saturation Levels

We expect the camera team to provide, for each amplifier, the lowest charge level at which charge moves from one pixel to a neighbour. If necessary, they will also provide the level at which the serial register saturates (i.e., if the serial saturates at a lower level than the parallels). Note that this is not the "saturation level" that the EOTest suite currently measures.

Take a set of exposures using the CBP with a mask with small holes to produce spots on the focal plane with peak flux levels at a set of linearly spaced intensity levels starting at c. 50000 DN and increasing by c. 1000 DN (i.e., 50000, 51000, 52000, ...). The exact levels are not important. When we know the devices' saturation levels we may be able to use a smaller set of levels. There should be a spot projected on every amplifier. If convenient, the parallels on the CCD may be run to transfer enough rows that the successive spots do not overlap to reduce the total data volume.

The individual frames will be processed through overscan, bias, dark, non-linearity, flatfield, brighter-fatter, and CTE (if available) ISR processing.

Tests:

1. Fit the PSF to the spots in the faintest image (either as a function of position or per-spot, depending on the properties of the mask). Measure the aperture and PSF fluxes of each spot. For each amplifier, plot the ratio as a function of peak intensity, fit a line to the faint end (correcting for errors in the B-F correction and non-linearity curve) and determine the point at which the ratio deviates from the fit.
2. Using the known values of the amplifier gains, confirm that the saturation level in electrons is consistent for all amplifiers in any one detector. N.b. this tests both the saturation levels and the gain determinations.
3. For spots above saturation (as determined above) measure the value of the bleed-trail as a function of distance from the centre of the spot. Allowing for spill-over into the neighbouring pixel at the top and bottom of the trail, find the lowest level present in bleed trails for each amplifier. Compare with the saturation levels for each amplifier determined above.
4. Repeat this analysis on real saturated star images when available.
5. Examine the levels of pixels to the left and right of bleed trails (i.e., in parallel to the serial register) for signs of saturation in the serial register.

9 Crosstalk

We expect the camera team to deliver a crosstalk matrix coupling every amplifier to every other amplifier; we hope and expect that most of these matrix elements will be 0. It is currently expected that crosstalk between separate CCDs will be unmeasurably small.

Using the CBP with a mask with reasonably large holes, produce spots on the focal plane with intensities well within the linearisable range (c. 50000DN). The number of spots in the mask depends on the degree of crosstalk present; e.g. just one; one or a few per CCD; one per amplifier. The experiment must be repeated until a spot has been placed upon every amplifier in the camera.

For a selected set of amplifiers, take the spot data at a set of logarithmically spaced intensity levels starting at 500 DN and increasing by a factor of 2 (i.e., 500, 1000, 2000, ...). The exact levels are not important. Depending on the spot areas, and cross-talk coefficients, it may be necessary to take multiple exposures at the same intensity level to measure the coefficients

with an accuracy sufficient to enable cross-talk correction without adding noise.

Take an exposure with a c. 50000DN spot projected onto every amplifier.

The frames will be processed through overscan, bias, dark, non-linearity, and flatfield ISR processing. Cross-talk correction should be included for the final, all-amplifier-spot, image.

Tests:

1. Fit the amplitude of each cross-talk image to the image of the "source" spot, resulting in a 3024x3024 matrix (16*189). If appropriate, a 16x16 matrix may be sufficient. Compare with the numbers provided by the camera team.
2. Using the all-amp-spots image, mask the direct spots, and confirm that the median and 5-sigma clipped means are equal to within statistical error.
 - If the scattered/ghost light from the CBP has too much structure this will fail and more sophisticated analysis will be needed.
 - Note that this requires the area covered by cross-talk artefacts to be small; if this is not satisfied, we will need to take a set of images with fewer spots.

We can also check the cross talk coefficients using saturated stars' bleed trails or cosmic rays (Yagi, 2012).

10 WRITE ME Flat Fields

Even with high-QE devices such as those in the LSST camera, the sensitivity of the detectors is a function of the SED of the light illuminating them (especially in the uv and above c. 950nm); this is made worse by any spatial structure in the transmission of the filters or of any other optical elements in the system. Additionally, the light received at a point in the focal plane is the sum of the focused light and of any ghost or scattered contribution.

For photometry of objects only the focused part is of interest, while for background subtraction it is the total incident light that matters.

We are considering applying corrections to the data for the static part of the pixel size non-uniformity. It is *very* important that the same corrections be applied to the flat fields.

10.1 Broadband and Monochromatic Flats

We will take both broadband *broadFlat* and *monochromatic* flat fields. As discussed below, we will use the latter to estimate the correct flat field for light with any desired SED (including that of the night sky), and the former to correct our data for the effects of dust on the filters and other effects that change on a short time scale.

10.2 Contamination of Filters and Other Optical Elements

Dust motes appearing on filters have an affect upon the system flat fields. This is not expected to be a serious effect for LSST as the beam at the filter has a diameter of c. 10cm (effective diameter c. 8cm) so only an unexpectedly large contaminant will have a measureable effect.

However, we shall take broadband flats every time we put a filter into the filter changer, and probably as a routine part of afternoon checkout. We will synthesise a broad-band flat from the monochromatic flats and filter transmissions, and divide the measured by the synthetic flats; by taking contaminants to be gray (e.g. opaque, but with a small covering factor) we will use the ratio to correct the monochromatic flats for any changes since they were taken.

Another approach is to use a broadband flat taken at the same time as the monochromatic flats as the reference image. We will do both; any discrepancy will be a sign of potential drift in the filter curves.

10.3 Flats for Background Estimation

The calibration system will produce a low-resolution spectrum of the night sky, and we will use this to synthesise a flat field image for the sky's SED from the monochromatic dome flats.

We expect that this image will have large scale structures different from those seen in sky image that are due to illumination gradients on the flat field screen, the screen's non-uniform BDRF, and the fact that the screen is not at infinity. We do not expect these gradients to cause problems in the background subtraction, and will not discuss them further except to say that

if desired we could correct them using median night-time sky frames.

Note that we are not proposing using twilight or nighttime sky flats as they do not have the desired SED and may cause extra problems due to polarization.

10.4 Flats for Photometry

The flat fields described in Sec. 10.3 are not suitable for photometry of astronomical objects for at least two reasons; they include indirect light, and they have the wrong SED. We can use the Collimated Beam Projector (CBP; CBP) in conjunction with the monochromatic flats `monoFlat` and filter transmission curves `filterTransmission` to solve this.

All the CBP data will be processed using the standard LSST ISR, except that no flat fields will be applied. We will then use standard LSST aperture photometry to measure the number of counts in each CBP spot.

- Estimation of the Photometric Flats at a Finite Number of Points.

If all the spots projected by the CBP were known to have the same intensity, then the spot fluxes measured from a scan in wavelength `CBP:mono` would give us the relative QE at a set of points in the camera. In reality the spots are not all equal (due to an imperfect mask, non-uniform illumination, and varying plate scale) so we need to solve for their relative intensities; we can do this by moving the spots around the camera, `CBP:spot`, in a manner similar to the standard processing of star flats.

We can then correct the monochromatic spots `CBP:mono` for the gains `gain`, and arrive at the relative QE at a set of points in the camera, as a function of wavelength, in the absence of a filter.

- Estimation of the Photometric Flats for All Pixels

These spot data `CBP:mono` only sample one point on M1, and therefore only one path through the optical system. We then use the data taken at different points in M1 `CBP:M1` to correct these data so that they reflect the performance of the entire optical system. We could sample all of M1, but in practice we expect to use many fewer pointings.

If we could repeat this operation putting the CBP spots down on every pixel we would have our desired flat field; unfortunately this is impracticable.

Instead, we will take the monochromatic dome flats `monoFlat` and use the known QE at the position of the CBP spots (without applying the filter transmission curves) to correct for ghosting. A sketch of the procedure is:

- Fit a surface through the CBP values (either per-CCD or for the whole camera); a spline would be a reasonable choice (either the product of two 1-D splines, or a thin plate spline. I would start with the former as they are easier to understand).
- Divide the dome flat by this surface, giving an estimate of the illumination and chip-to-chip correction
- Fit a curve to this correction, and correct the dome screen. This should be close to the values derived from the CBP data (and can preserve discontinuities in the QE across chips which the fitted curves have a hard time following).
- Iterate a couple of times; each iteration should result in a smaller and smoother correction, which we are therefore better able to model.

Repeat this operation at a suitable set of wavelengths, chosen so that the variation of these corrections as a function of wavelength is well captured; we now know the relative QE for all pixels in the camera, as a function of wavelength, in the absence of a filter.

Using the filter transmission curves `filterTransmission` we can derive the relative QE for all pixels in the camera for each filter at 1nm resolution; this is our monochromatic photometric flatfield.

N.b. we are in a position to provide a model of the ghosting following this analysis in two ways:

- By analysis of the CBP data, concentrating not on the spots but on the scattered/ghost light
- By looking at the corrections applied to the CBP spot data to arrive at the dome screen data

We do not need this information to calibrate LSST, but it will provide a valuable cross-check, and will inform the camera and telescope teams about the state of the optical elements.

10.5 Outputs

The corrected monochromatic flats (appropriate for the background and photometry) will be the inputs to pipeline processing; `monoFlatmonoPhotoFlat`.

We will also employ a photometric flat appropriate to a flat-spectrum source (or other well-defined SED) `monoPhotoFlat` absorbed by a standard atmosphere. This is not strictly-speaking an output — it could be constructed on the fly — but it is more efficient to construct it here. `standardPhotoFlat`. *N.b.* the use of a standard SED/atmosphere here is intended to make the corrections for objects' real SEDs/atmospheric absorption smaller and with smoother spatial structure. Its use is not intrinsic to the final photometry, and if it turns out not to be helpful it may be dispensed with.

If, as we expect, the error due to using this standard spectrum flat is slowly varying and only dependent on a low-resolution version of the SED (*e.g.* a resolution sufficient to capture the photometric effects of an Ia SNe), we will prepare a low-resolution data cube to capture the essence of `monoPhotoFlat`; `monoPhotoFlatLowRes`.

11 WRITE ME Fringe Correction

Although the thick, deep-depletion LSST CCDs reduce the amplitude of fringing they do not remove the necessity to handle it correctly in the z and y bands. Fringing is a QE effect, and this QE variation coupled to night sky line emission (mostly from OH) leads to spatial structures in the background.

The classical approach to fringing is to subtract a multiple of a fringe frame, made by taking the median stack of a very large number of science exposures scaled by their sky intensity. Because the lines in the night sky spectrum vary in relative intensity during the night the structure of the fringing will, in general, vary.

In addition, the night sky shows spatial structures in rotational temperature and intensity with wavelengths of *c.* 2km (the OH emission is at *c.* 100 km, so wavelengths of *c.* 1 degree, or 4 CCDs). With a windspeed of 60 km/s the smoothing scale (in the wind direction) is *c.* 1 CCD. These spatial structures change on time scales of *c.* 10 minutes.

If a CCD is of more-or-less of constant thickness with small ($\ll 1\mu\text{m}$) variations the relative variation of the line intensities do not significantly affect the structure of the flats as the fringe patterns from different lines are very similar (although of different amplitude and phase). However, at least some of the LSST chips show 20 fringes at $1\mu\text{m}$ from centre to edge of the device, which implies that the *pattern* of the fringing will change as the line ratios change.

For small fields of view, or long integrations, where the spatial structure in the night sky brightness may be ignored it is possible to model the variation in the fringing using a PCA decomposition of sky frames. Unfortunately LSST will show changes in fringe patterns and intensities coming from both the spectral and spatial variation of the sky and we are planning to use a different approach.

As described in Sec. 10.3 we will synthesise flat fields that match the night sky's SED. In a little more detail relevant to fringing, the calibration system will produce a spectrum of the night sky with resolution $c. 200$, and this is sufficient to predict the distribution of night sky line intensities `nightSkySpectrum`. We will then use this low-resolution spectrum to construct a flat field from the monochromatic dome flats `monoFlat` that flatten the night sky spectrum, removing small-scale structure due to the sky emission.

Note that we assume that a single spectrum within the field of view is sufficient to constrain the night sky spectrum, whereas in reality the (atmospheric!) gravitational waves producing spatial structure have the ability to modify the emission spectrum. As timescales in the atmosphere are $c. 10-20$ minutes we will have many realisations of the spectrum at different points in the sky, and will be able to use these to confirm that the effective fringe image is spatially uniform; if it transpires that this is not the case we will be able to construct a small number of flats that capture the variation, and use them to remove the variation in the spatial structure.

If this sky-flattening approach works we will not need any special fringe outputs; we should allow for the use of some variant of fringe frames as a backup `fringeFrames`.

12 WRITE ME Filter Transmission Curves

The filter transmission curves `filterTransmission` will be provided by the camera team, but by using monochromatic spots from the CBP with `CBP:filterCBP:leak` and without filters in the beam `CBP:mono` we can monitor the filters for any evolution, including light leaks.

Note that we are not exploring all possible light paths through the filters; while 2.5 CBP spots per CCD mean that we have light passing through every point on the filter they are not passing through every patch on the filter at every angle.

In theory we can measure the filter curves using the CBP, but it'd take a long time; in full

generality *c.* 470 exposures for each 1 nm step. If we are willing to make strong enough symmetry assumptions this can be reduced to *c.* 6 exposures per wavelength step; an extension of CBP:M1 to higher spectral resolution, and with the filter in the beam.

13 WRITE ME Pixel Size Effects

It has become clear that much of the pixel-to-pixel variation seen in flat field images is in fact due to variations in the sizes of the pixels. These are popularly divided into two parts:

- Large scale features such as the “tree-rings” due to variation in the properties of the silicon, or perturbations in the electric field (*e.g.* “edge distortions” [also known as “glowing edges” in certain devices]).
- Small-scale, quasi-random, variations in the pixel sizes.

If we interpret the flat fields in terms of QE variation we will make errors measuring fluxes (although not while estimating surface brightnesses).

The pixel size variation may be thought of as a 2-D vector field ξ , the offset of each corner of each pixel from a regular grid. There are various approaches being taken to measure the sizes of the pixels.

If you can reduce the distortion to a 1-dimensional function (*e.g.* tree rings with circular symmetry about a point; ‘glowing edges’ that are a function of distance from the edge of the sensor) then high-signal-to-noise flat fields can be used to determine the distortions broad-Flat.

The small-scale variations are harder. If we are willing to make the arbitrary assumption that ξ is the gradient of a scalar we can solve for it from flat field data. Aaron Roodman’s group, working on DECam data, makes the surprising claim that they can find ξ by looking for a solution close to $\xi = 0$.

If such approaches fail, it may be possible to measure ξ by using images of 1-D sinusoids projected onto the CCD, either by the camera team or on the mountain using the CBP.

See Sec. 13 for a discussion of how we will use ξ once it is known; `pixelSizeMap`.

14 WRITE ME Brighter-Fatter

The measurements needed to characterise Brighter-Fatter effects are not yet clear. Current state of the art is to use pixel-to-pixel correlations in flat fields taken at different flux levels `broadFlat`. It is possible that we'll need other measurements, possibly generated by projecting masks onto the camera using the CBP; `brighterFatterCoeffs`.

15 WRITE ME Ghosts and Ghouls

Write me

16 Photon Transfer Curve

The Photon Transfer Curve (PTC) of a CCD plots the variance of a series of images taken under uniform illumination ("flat fields") as a function of their average. It allows for the calculation of fundamental CCD parameters such as the gain (electrons per ADU), the readout noise, and the mean pixel full well via the PTC turnoff. As such, it plays an important role in the characterization and monitoring of the state of a detector.

It has been established that the variance flattens out at high fluxes and it is not perfectly proportional to the average flux, as Poissonian statistics would dictate. This is related to the "Brighter-fatter" (BFE) effect and increases the correlations between adjacent pixels, decaying with distance, as reported in e.g., Astier et al. 2019. This study proposes a couple of models to fit to the PTC data (including covariances) with a leading order parameter α_0 to take the BFE into account. This parameter should be negative (as it describes the change in a pixel area due to its own charge content), is purely electrostatic in nature, and should not change during the life of a given CCD unless operating voltages such as the bias voltage are changed.

Tests:

1. Compare the fitted gain to the gain value stored in each amplifier, to monitor any temporal changes. Check that the relative fractional difference is within a threshold (default: 5%).
2. Compare the fitted readout noise to the readout noise value stored in each amplifier, to monitor any temporal changes. Check that the relative fractional difference is within a threshold (default: 5%).
3. Check that the reported PTC turnoff value (converted to electrons) is greater or equal than the full-well requirement (90000 electrons).
4. Check that the brighter-fatter effect parameter from the Astier et al. 2019 models, a_{00} , is within a range determined by measurements of the distribution of this parameter (e.g., within five sigma of the mean of the coefficientis distribution for a given instrument, detector, and fit type).

A References

References

[LSE-61], Dubois-Felsmann, G., Jenness, T., 2019, Data Management System (DMS) Requirements, URL <https://lse-61.lsst.io/>,
Vera C. Rubin Observatory LSE-61

Yagi, M., 2012, PASP, 124, 1347 (arXiv:1210.8212), doi:10.1086/668891, ADS Link