DESI Commissioning Instrument CCD Characterization

1 Background Information

The Dark Energy Spectroscopic Instrument (DESI) Commissioning Instrument (CI) was an imager used to preliminarily verify key functionality of the Mayall telescope at Kitt Peak National Observatory as outfitted with its newly installed DESI corrector, including the tracking, guiding, pointing model, polar axis alignment, hexapod motion, throughput, focusing, and image quality. The CI was equipped with five commercial SBIG STXL-6303e cameras and 22 illuminated fiducials, and matched the mass and moment of the full DESI focal plane system that will be used throughout the upcoming five-year spectroscopic survey.

On-sky observing with the CI began on the evening of 2019 April 1 following CI installation at Kitt Peak the week prior. This marked the first on-sky observing with the Mayall telescope since 2018 February. The final on-sky CI observing occurred during the morning of 2019 June 3. Thus, the nights of on-sky CI campaign data acquisition spanned from 20190401 to 20190602 (inclusive), where these labels correspond to the calendar date (Kitt Peak local time) at the start of the night. There was an approximately 1.5 week 'pause' in the CI on-sky observing which lasted from 20190506 to 20190515 (inclusive), to move the corrector in z by approximately 9 mm and thereby center the best focus very close to the hexapod neutral z position.

Each CI camera contains a 3072 pixel $\times$ 2048 pixel detector, with square pixels 9μm on a side. Although all CI pixels are physically square, they can subtend non-square regions on-sky due to the directionally-dependent DESI corrector platescale, and also subtend different solid angles as a function of radius from the field of view center. The approximate CI platescale is $\sim$0.13$''$/pixel. Figure 1 shows a schematic of the CI camera layout within the focal plane (see DESI-3347 for more such details). Critically, one of the CI cameras ("CIC") is in the center of the focal plane, and therefore has provided the only opportunity for on-axis imaging through the DESI corrector. The outlying four CI cameras ("CIE", "CIN", "CIS" and "CIW") are separated from CIC along the sky E, N, S, W cardinal directions by an on-sky angular radius of $\sim$94 arcminutes. These outer four CI cameras are meant to help understand the astrometry, vignetting, image quality, and so forth at the very edge of the wide DESI field of view, and to mimic the guide-focus-alignment cameras (GFAs) included in the DESI focal plate assembly, although the GFAs will not be SBIG STXL-6303e devices like those of the CI. The CI cameras all observe through a narrow Astrodon r' filter transmitting in the wavelength range 560 nm $\lesssim \lambda \lesssim$ 700 nm[1].

Other information closely related to that presented in this document can be found in Ross et al. (2019; DESI-3624), DESI-3516, DESI-3358, DESI-3347, DESI-4790, and the CI CCD data sheet[2].

2 Camera Labeling

One important cautionary note to keep in mind is that prior to 20190402, four out of the five CI camera labels were "swapped" within the raw metadata relative their sky locations. Throughout this document, we refer to each camera by its sky location[3]. For data up to and including 20190401, the camera labeled "CIE" was actually CIW, the camera labeled "CIW" was actually CIE, the camera labeled "CIN" was actually CIC, and the camera labeled "CIC" was actually CIN. This label swapping is summarized in Table 1. CIS was always labeled correctly. Starting on 20190402, the labels in the raw metadata were corrected so that they matched the sky locations of all five cameras.

3 Calibration Data

The CI observing instructions did not stipulate any formal requirements about collecting specific/standardized calibration data sets on e.g., a nightly basis during the CI run. Nevertheless, a substantial amount of calibration data

[1] https://astrodon.com/products/astrodon-photometrics-sloan-filters/

[2] https://www.onsemi.com/pub/Collateral/KAF-6303-D.PDF

[3] The four outlying cameras would be labeled differently if the labels were based on geographic location, so it is important to clarify that our labels refer to sky location.

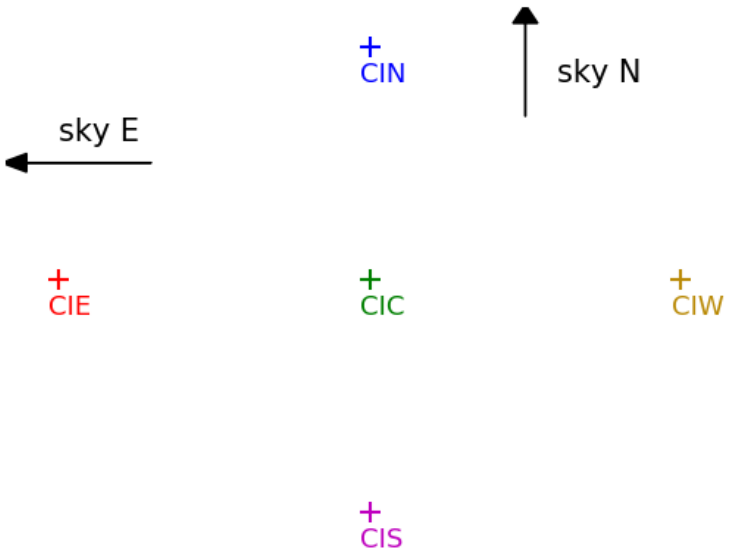

Figure 1: CI naming convention used in this work. Each camera is referred to by its sky location. The names of the four outlying CI cameras would be different if names based on their geographic locations were adopted. The names of all but CIS would be different if the "swapped" raw metadata labels from 20190401 and earlier were adopted.

label through 20190401	adopted camera name (sky location)
CIE	CIW
CIN	CIC
CIC	CIN
CIS	CIS
CIW	CIE

Table 1: Correspondence between camera name adopted throughout this work (based on sky location, right column) and "swapped" camera labels in the very early portion of the CI run (night 20190401 and earlier). CIS was correctly labeled in the raw metadata even on 20190401.

were collected sporadically throughout the CI campaign. Before a full calibration analysis can be undertaken, a precursor effort is needed to identify and catalog all relevant data sets gathered throughout the entire CI campaign.

There are two primary methods for most conveniently identifying CI data sets of interest for a particular analysis task. The first is a parameter called FLAVOR, and the second is a parameter called PROGRAM. The FLAVOR keyword directly specifies the general type of data being acquired. Examples include FLAVOR=zero (or FLAVOR=ZERO) for bias frames, FLAVOR=dark (or FLAVOR=DARK) for dark frames, and FLAVOR=science (or FLAVOR=SCIENCE) for exposures where the CI camera shutters are opened to collect light. The FLAVOR keyword is not meant to be granular. For instance, there is no distinct FLAVOR value for flat field images, since these are simply one of many possible types of data that fall under the FLAVOR=science umbrella. The PROGRAM keyword provides the potential for additional granularity. Each CI observation request contains a PROGRAM keyword which observers could use to communicate in detail the intended purpose of the data being acquired. Thus, although there is no FLAVOR that specifically identifies flat field images, one would hope that strings such as 'flat' and/or 'twilight' were input in some way into the observation request PROGRAM values by the observers when acquiring twilight flats. The PROGRAM and FLAVOR keywords are available, for example, in the raw CI image header metadata.

In the following subsections I describe my "inventories" of available post-installation CI bias (§3.1), dark (§3.2) and flat field (§3.3.2) images based on analyzing the available FLAVOR and PROGRAM metadata. Very minimal pre-installation CI engineering data were made available – a small handful of bias frames sufficient to measure the read noise and conceivably construct noisy "master" bias images, but no data which could be used to accurately measure any of the gain, detector flat fields, or dark current. Furthermore, given the label swapping discovered

during the first night of on-sky CI observing (§2), it would remain unclear whether we should trust pre-installation engineering to have reliable camera labels. Taken together, these are the reasons why we prefer to and/or must use post-installation CI data rather than pre-installation engineering data throughout this study.

3.1 Inventory of Available Bias Sequences

I gathered the metadata for all raw ci*.fits.fz files on disk at NERSC in night directories with NIGHT $\geq$ 20190330. I then restricted to the subset of exposures with flavor=zero (any capitalization), and partitioned these into "sets" by demanding that a "set" consist of consecutive EXPID values[4] and not jump between multiple night directories. This led to 91 distinct sets of biases. Table 2 has one row per set of biases with at least five consecutive exposures and is sorted by EXPID. The rightmost column labels whether the data have the CIN/CIC and CIE/CIW camera label swapping issue described in §2. This flag is just based on night $\geq$ 20190402 — if early calibration exposures on 20190402 may possibly have their camera labels still swapped then it would be safest to discard bias sets from 20190402. I verified that all bias frames contributing to this inventory have EXPTIME = 0, as expected (EXPTIME is another CI metadata keyword which provides the exposure time in seconds).

31 of these sets of biases have at least 5 exposures. 19 of these sets of biases have at least 7 exposures. 16 of these sets have at least 9 exposures. A more extensive version of Table 2 plus the same data in FITS binary format are available on the DESI wiki[5].

3.2 Inventory of Available Dark Sequences

I analyzed the metadata for all raw ci*.fits.fz files on disk at NERSC in night directories with NIGHT $\geq$ 20190330. I then restricted to the subset of exposures with flavor=dark (any capitalization), and again partitioned these into "sets" by demanding that a set consist of consecutive EXPID values and not jump between multiple night directories. This led to 28 distinct sets of darks. Table 3 has one row per set of darks and is sorted by EXPID.

17 of these sets of darks have at least 5 exposures. 13 of these sets of darks have at least 7 exposures. 12 of these sets have at least 9 exposures. A more extensive version of Table 3 plus the same data in FITS binary format are available on the DESI wiki[6].

3.3 Flats

There are a variety of ways one might imagine attempting to characterize the CI flat field. Ideally, a setup similar to that used recently for MzLS at the Mayall would be available, with dedicated calibrations lamps uniformly illuminating the Mayall/DESI calibration screen. During MzLS, this setup allowed us to obtain a standard sequence of dome flats every afternoon as part of the observing procedure[7]. Unfortunately, the calibration lamp setup previously available at the Mayall has been removed as part of DESI installation, and the new DESI calibration lamp assembly has not yet been installed. As a result, for the CI run, we had no means of illuminating the Mayall/DESI calibration screen in a spatially uniform way, as would have been desirable for mapping the CI flat field.

Additionally, no CI detector flats were obtained in the laboratory prior to CI installation, making it challenging to decompose any CI flat field measured post-installation into the separate contributions from the detector response versus the telescope/corrector.

On 20190402, we attempted to take dome flats in evening twilight by imaging the calibration screen illuminated by ambient light with the dome open (EXPID = 3442-3445). However, the illumination was horribly non-uniform and clearly not suitable for deriving flat field corrections. At various times we took calibration screen data with the dome closed and the "low dome" lights providing ambient illumination (§3.4). While these low dome light calibration screen data ultimately proved very valuable (see §8), the calibration screen illumination was still highly non-uniform and therefore unsuitable for deriving flat field corrections.

This left only one potentially viable option for attempting to acquire flat field calibration images: twilight flats.

[4] EXPID is an integer counter that uniquely identifies each CI exposure and increases with time.
[5] https://desi.lbl.gov/trac/wiki/Commissioning/CI/bias_inventory
[6] https://desi.lbl.gov/trac/wiki/Commissioning/CI/dark_inventory
[7] https://desi.lbl.gov/trac/wiki/PublicPages/MayallZbandLegacy/NotesforObservers/Domeflats

first EXPID	last EXPID	# consecutive exposures	night directory	camera label swapping?
2727	2737	11	20190330	1
2796	2895	100	20190330	1
3107	3111	5	20190331	1
3196	3248	53	20190401	1
3407	3417	11	20190402	0
3437	3441	5	20190402	0
3446	3450	5	20190402	0
3685	3689	5	20190403	0
4174	4194	21	20190406	0
4204	4211	8	20190406	0
4218	4227	10	20190406	0
6246	6250	5	20190412	0
7087	7102	16	20190414	0
7635	7641	7	20190417	0
7652	7657	6	20190418	0
8863	8867	5	20190422	0
11218	11280	63	20190504	0
11592	11655	64	20190505	0
11983	11987	5	20190516	0
12336	12340	5	20190517	0
12643	12647	5	20190518	0
12766	12777	12	20190519	0
12831	12835	5	20190520	0
12945	12951	7	20190521	0
12956	12960	5	20190522	0
13313	13322	10	20190526	0
13433	13532	100	20190526	0
13651	13750	100	20190526	0
13853	13868	16	20190527	0
14063	14079	17	20190528	0
14091	14105	15	20190528	0

Table 2: Inventory of all bias sequences with at least 5 consecutive bias frames, from 20190330 onward, sorted by exposure number (EXPID). There are 31 such bias sequences in total.

3.3.1 Twilight Flat Observing Procedure

To acquire sky flats, we pointed near zenith during evening twilight[8] with the dome and mirror covers open and took a series of CI flavor=science exposures of increasing exposure time. We started off taking a rapid sequence of short 10 second exposures until those became unsaturated, then gradually increased the exposure time as the sky became darker until reaching exposure times of up to 600 seconds. Most of the twilight flat sequences during the CI run were taken by me, and I used a real-time data analysis module that I implemented called `ci_imstats` to monitor the median count level of the most recently acquired twilight flat exposure in each camera and adjust the subsequent exposure time(s) accordingly[9]. The goal was to remain unsaturated while accumulating at least 10,000 ADU of background per exposure, with a background level of several tens of thousands of ADU being considered optimal. Between consecutive twilight flat exposures we dithered by manually sending target offset requests via the DELTARA, DELTADEC ICS keywords, with the dither scale being several tens of arcseconds. Dithering avoids having the same stars land on the same detector pixels in all/many frames within a given twilight flat sequence. Sometimes twilight flats were taken with the telescope roughly in focus and sometimes we were far out of focus. But this shouldn't have mattered provided we did the usual dithering. The dithers were entered manually rather than

[8] Twilight flats were never taken in morning twilight during the course of the CI run.

[9] `https://github.com/ameisner/ci_reduce/blob/master/py/scripts/run_imstats.py`

first EXPID	last EXPID	# consecutive exposures	night directory	camera label swapping?	# unique EXPTIME	min EXPTIME	max EXPTIME	total EXPTIME
2738	2747	10	20190330	1	1	2	2	20
2905	2974	70	20190330	1	1	300	300	21000
3094	3102	9	20190331	1	1	300	300	2700
4094	4098	5	20190405	0	1	15	15	75
4142	4153	12	20190405	0	1	5	5	60
6204	6240	37	20190411	0	1	300	300	11100
10702	10714	13	20190502	0	1	300	300	3900
11988	11992	5	20190516	0	1	3	3	15
12341	12350	10	20190517	0	1	300	300	3000
12778	12787	10	20190519	0	1	300	300	3000
12961	12965	5	20190522	0	5	1	16	31
13008	13012	5	20190523	0	5	1	16	31
13053	13065	13	20190524	0	8	1	128	286
13183	13190	8	20190525	0	8	1	128	255
13323	13432	110	20190526	0	2	5	10	1050
13751	13850	100	20190526	0	1	10	10	1000
14080	14090	11	20190528	0	1	300	300	3300

Table 3: Inventory of all dark sequences with at least 5 consecutive dark frames, from 20190330 onward, sorted by exposure number (EXPID). There are 17 such sequences in total.

scripted and so in a few instances they were accidentally forgotten. Typically we were unfortunately not paying attention to whether the ADCs were in neutral, nor were we trying to systematically rotate the ADCs to search for the differences in vignetting that would be predicted to result.

3.3.2 Inventory of Twilight Flat Sequences

I gathered the metadata for all raw ci*.fits.fz files on disk at NERSC in night directories with NIGHT $\geq$ 20190330 and with flavor=science (any capitalization). I then restricted to those exposures that had one or more of "twi" or "flat" (any capitalization) in their PROGRAM string, and subsequently removed any exposures with any of "dome", "shutter" or "conditions" in their program string (any capitalization). This left 179 twilight flat exposures. I then partitioned these 179 exposures into sets by demanding that a set consist of consecutive EXPID values and not jump between multiple night directories. This led to 12 distinct sets of twilight flats from 11 different nights. Table 4 has one row per set of twilight flats and is sorted by EXPID.

11 of these sets of twilight flats have at least 5 exposures. 10 of these sets of twilight flats have at least 7 exposures. 9 of these sets of twilight flats have at least 9 exposures.

3.4 Calibration Screen Data

This test should be done with the dome closed at night.

- ask the OA to point the telescope at the calibration screen
- turn the low dome lights on (but not any other lights)
- mirror covers need to be open
- tracking and dome movement need to be disabled

The actual observations can be taken using a script generated by the Python code at:
`https://github.com/desihub/desici/blob/master/py/desici/gain_low_dome.py`

first EXPID	last EXPID	# consecutive exposures	night directory	camera label swapping	# unique EXPTIME	min EXPTIME	max EXPTIME	total EXPTIME
4543	4554	12	20190407	0	6	5	600	1160
4932	4944	13	20190408	0	9	10	400	960
5748	5754	7	20190410	0	7	10	600	1170
6550	6561	12	20190413	0	7	2	30	148
7415	7423	9	20190417	0	6	20	120	515
8317	8347	31	20190420	0	13	2	140	778
8524	8535	12	20190421	0	9	10	200	582
10954	10958	5	20190503	0	1	10	10	50
10970	10994	25	20190503	0	8	10	500	1215
13870	13900	31	20190527	0	11	1	600	1712
14108	14128	21	20190528	0	8	10	400	1160

Table 4: Inventory of all twilight flat sequences with at least 5 consecutive dark frames, from 20190330 onward, sorted by exposure number (EXPID). There are 11 such sequences in total.

This scripting only deals with taking the CI images. The setup steps (e.g., having the OA point the telescope to the calibration screen and turn on the low dome lights) should be done manually. The observations gradually increase in exposure time from 2 seconds to 32 seconds in steps of 2 seconds, with (by default) 4 observations per unique exposure time, and all observations at a given EXPTIME acquired in a contiguous block. The idea is to measure the count level and variance at each unique EXPTIME value, with the variance measured by taking pairwise image differences. There is no need for the illumination to be spatially uniform (i.e., these aren't meant to be flat field images). However, it is very important that the overall illumination level not vary substantially on the timescale of the 4-exposure sequence taken at each EXPTIME value. Based on the 20190411 data set, the longest exposures in this sequence should be saturated or close to saturated in all five cameras. When we did this test early in the CI campaign, we ran into some challenges taking CI "science" exposures with the dome closed, such as hitting interlocks, so it may be best to try this test early in the night when an expert like Behzad or Klaus is available to help in case such situations occur again.

The sequence of CI exposures themselves should take 25 minutes assuming 4 frames per unique EXPTIME. The setup time could be a significant overhead depending on whether the ICS gets in the way of taking CI "science"images while pointed at the dome...

4 Bias Analysis

4.1 Bias Levels

For each exposure contributing to any of the bias 'sets' in Table 2, we computed the median bias level in each of its five cameras. We then grouped these median bias level measurements by set (i.e., by row of Table 2), and computed the median by camera of the per-exposure medians within each set. We thereby obtained a per-camera measurement of the overall bias level for each bias sequence. Figure 2 plots these bias level results as a function of time throughout the CI campaign, with each camera represented by a line of a different color. The changes in median bias level seem to be quite well correlated across four of the cameras: CIN, CIC, CIS, and CIW. CIE appears to show substantially less correlation with the other four cameras. One notable difference between CIE and the other cameras is that CIE's cooling was turned off/on at various times throughout the CI run (e.g., [desi-commiss 961]), although we have not proven that this is the underlying cause of the discrepant CIE behavior shown in this plot.

4.2 Master Bias Images

We created master bias images for each camera by combining the 100 bias frames from the sequence in the second row of Table 2 (EXPID = 2796-2895) using a per-pixel sigma-clipped mean. These master bias images (Figures 3, 4, 5, 6 and 7) are very cosmetically clean overall, with only a few discrepant columns visible per camera and show only very low-level large-scale spatial structure ($\lesssim$1%). As a result, subtracting the bias during raw image reduction

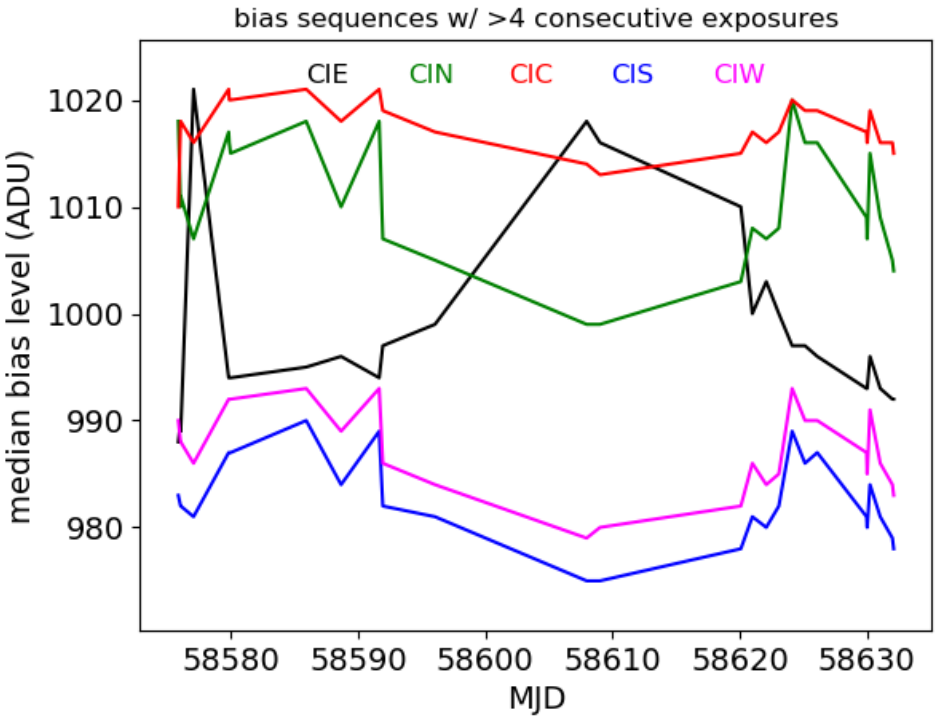

Figure 2: Median bias level for each of the five CI cameras as a function of time. Each data point represents the result of analyzing a single "set" of at least five consecutive bias exposures. The changes in median bias level seem to be quite well correlated across four of the cameras: CIN, CIC, CIS, and CIW. CIE appears to show substantially less correlation with the other cameras. One notable difference between CIE and the other cameras is that CIE's cooling was turned off/on at various times throughout the CI run (e.g., [desi-commiss 961]), although we have not proven that this is the underlying cause of the discrepant CIE behavior shown in this plot.

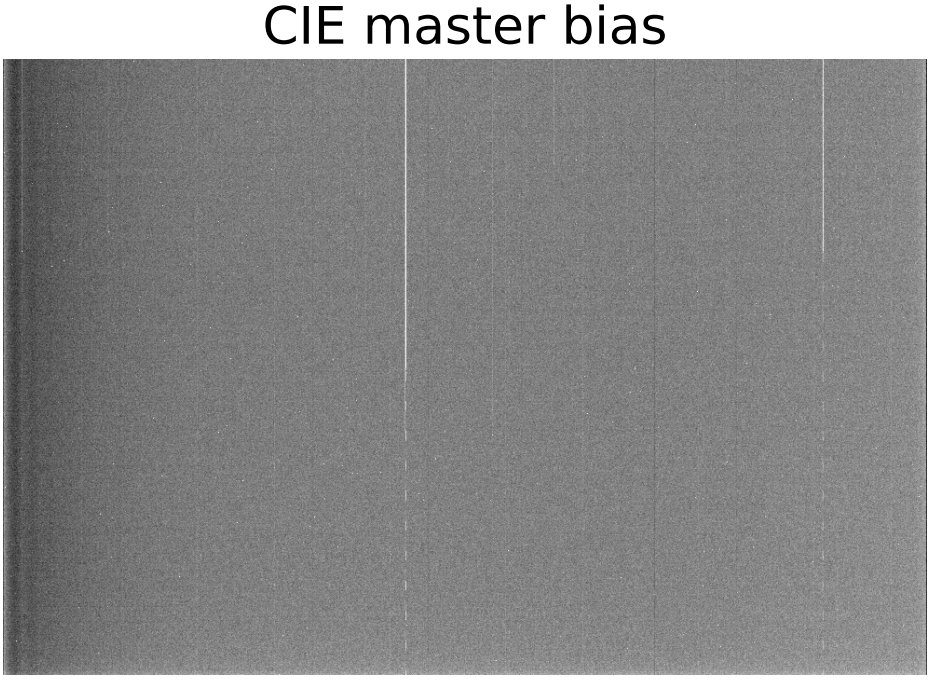

Figure 3: CIE master bias. The grayscale shown renders the ratio of the master bias to the median of this camera's master bias, with the stretch ranging from 0.99 (black) to 1.01 (white). Aside from a very small number of problematic columns, the spatial structure of the CI bias images is not particularly strong, and subtracting the bias during image reduction generally does not yield a strikingly obvious improvement in cosmetic appearance.

does not yield a strikingly obvious improvement in cosmetic appearance. Some of the relatively "hot" pixels in the bias frames are likely the result of pixels with unusually high dark current aliasing into the bias measurements due to the non-zero read-out time.

4.3 Readnoise

We group bias exposures into sets according to Table 2, and compute the difference between each pair of consecutive exposures within each set. Each pairwise difference yields a measurement of the readnoise (in ADU) based on computing the standard deviation of the difference (using a Gaussian model fit with outlier rejection) and dividing that standard deviation by $\sqrt{2}$. These per-pair readnoise estimates are averaged across all pairwise differences within each set to provide per-sequence readnoise measurements for each camera. These per-sequence readnoise measurements were then averaged together to provide the values quoted in Table 5. No strong time variation in the readnoise was observed.

camera	readnoise (ADU/pixel)	readnoise ($e-$/pixel)	assumed gain ($e-$/ADU)
CIE	8.35	14.49	1.735
CIN	8.27	14.10	1.705
CIC	8.40	14.37	1.710
CIS	8.14	13.72	1.685
CIW	8.30	14.52	1.750

Table 5: Summary of readnoise measurements, which result from combining information from all sets of bias sequences listed in Table 2. The readnoise measurements are performed in the native ADU/pixel units of the raw CI images (second from leftmost column), then converted to $e-$/pixel (second from rightmost column) via multiplication by the per-camera gain. The assumed gain values are taken from Table 7.

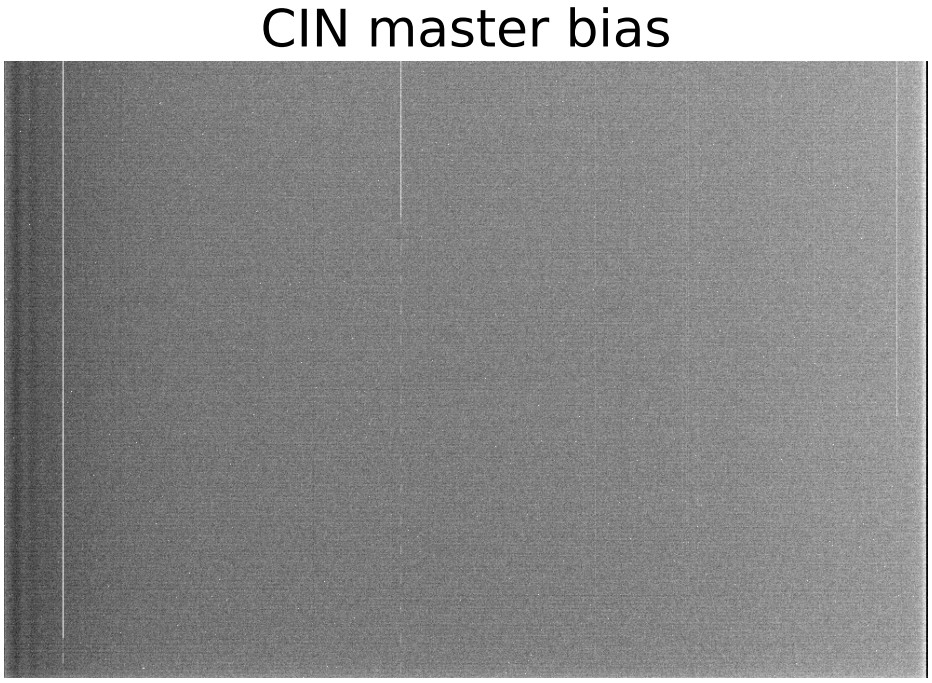

Figure 4: CIN master bias. The grayscale shown renders the ratio of the master bias to the median of this camera's master bias, with the stretch ranging from 0.99 (black) to 1.01 (white). Aside from a very small number of problematic columns, the spatial structure of the CI bias images is not particularly strong, and subtracting the bias during image reduction generally does not yield a strikingly obvious improvement in cosmetic appearance.

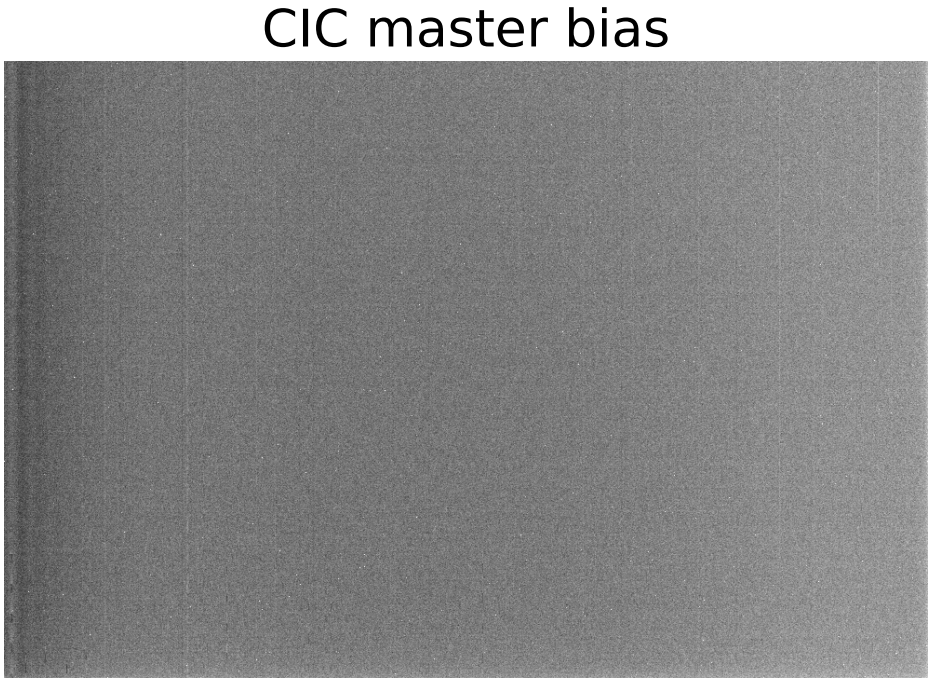

Figure 5: CIC master bias. The grayscale shown renders the ratio of the master bias to the median of this camera's master bias, with the stretch ranging from 0.99 (black) to 1.01 (white). Aside from a very small number of problematic columns, the spatial structure of the CI bias images is not particularly strong, and subtracting the bias during image reduction generally does not yield a strikingly obvious improvement in cosmetic appearance.

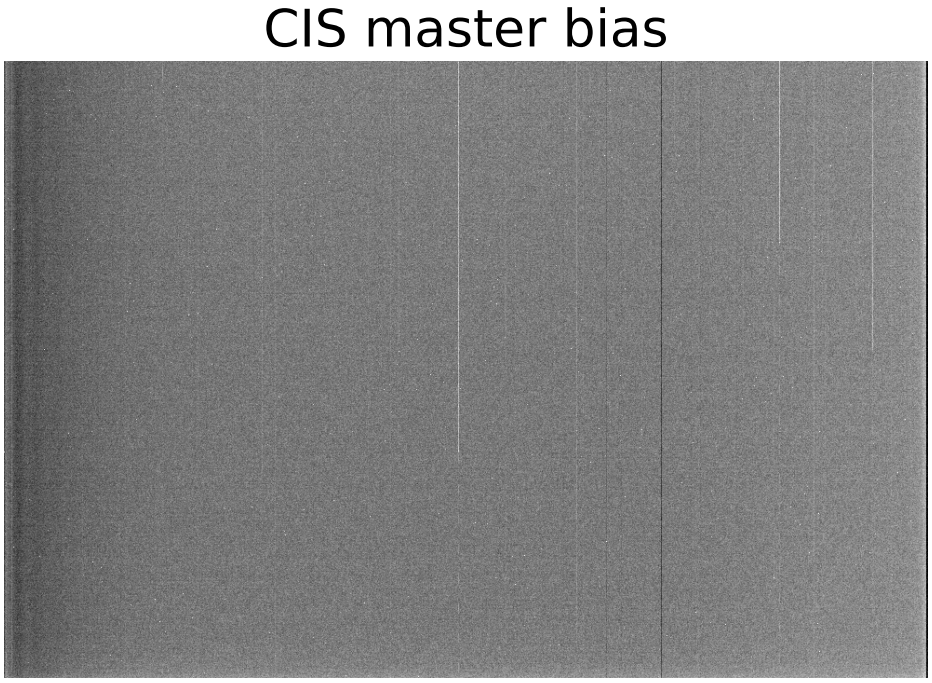

Figure 6: CIS master bias. The grayscale shown renders the ratio of the master bias to the median of this camera's master bias, with the stretch ranging from 0.99 (black) to 1.01 (white). Aside from a very small number of problematic columns, the spatial structure of the CI bias images is not particularly strong, and subtracting the bias during image reduction generally does not yield a strikingly obvious improvement in cosmetic appearance.

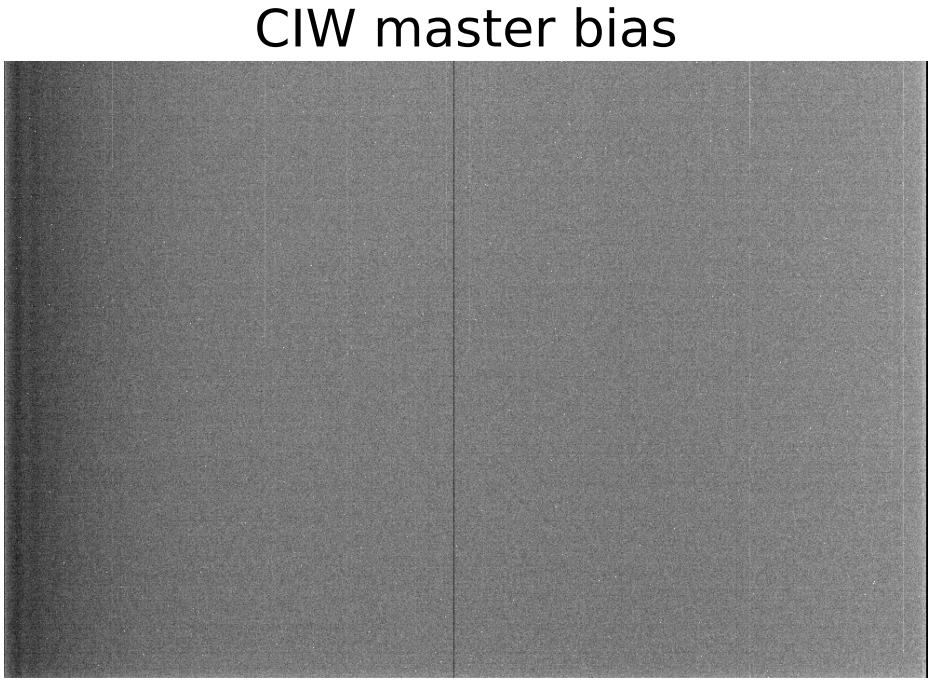

Figure 7: CIW master bias. The grayscale shown renders the ratio of the master bias to the median of this camera's master bias, with the stretch ranging from 0.99 (black) to 1.01 (white). Aside from a very small number of problematic columns, the spatial structure of the CI bias images is not particularly strong, and subtracting the bias during image reduction generally does not yield a strikingly obvious improvement in cosmetic appearance.

Readnoise values in $e-$/pixel in Table 5 assume the gain values in the gain measurement table that appears later in this document. These readnoise values take into account all bias sequences with $\geq$5 consecutive exposures, from 20190402 onward. Describe the methodology: pairwise differences of consecutive bias exposures within each set.

5 Flat Field Results

Unfortunately it can be seen (e.g., Figure 8) that the twilight flats are subject to significant contamination from light leaks. We therefore prefer to measure the flat field using star flats rather than flat field calibration images. Star flats are discussed and presented in a companion document focused on the CI throughput analysis.

6 Dark Current

6.1 Creation of Master Dark Images

To create the set of CI master dark images discussed throughout the remainder of this section, we performed a sigma-clipped mean combination of the set of 70 300 second dark frames acquired on 20190330 (see Table 3). We have correctly accounted for the camera label swapping that was in effect at that time. The master dark images[10] are natively in units of ADU per pixel per second. Conversion to $e-$/pix/second requires multiplication by an assumed gain value. Table 6 lists the median dark current rate value in each camera at the set point temperature of 7.5 C, since on 20190330 all cameras were run with cooling on. In data from that date no CCDTEMP information is available in the CI headers, and CCDTEMP information remained both inscrutable and partially absent throughout the CI run, hampering efforts to fully understand and correct for the dark current (see e.g., [desi-commiss 784], [desi-commiss 925]).

Averaging over the five CI cameras, the mean of the median dark current rates at 7.5 C from Table 6 is 0.25 $e-$/pix/sec. We can compare this measurement of the overall CI dark current rate to that from pre-installation acceptance analyses. In DESI-3358, CI dark current rates of 0.16 $e-$/pix/sec (0.38 $e-$/pix/sec) were reported at 5 C (10 C). Interpolating linearly between the dark current values at these two temperatures, (which admittedly isn't exactly the right way to do the interpolation) one would quote the DESI-3358 dark current rate at 7.5 C as (0.38 + 0.16)/2 $e-$/pix/sec = 0.27 $e-$/pix/sec, very close to our post-installation measured value.

6.2 Spatial Structure

The low-order structure of the master dark images is dominated by "tree ring" patterns for all cameras. This tree ring structure is of sufficiently low amplitude that it is not especially evident upon simply inspecting the native master dark images. Median filtering the master dark images makes the tree ring structure obvious, as shown in Figures 10, 11, 12, 13, and 14. Although the ring patterns are detected at high significance in these median filtered master dark images, their amplitude is quite small, almost always remaining within a few percent of the per-camera median value (see the Figure 10-14 captions for details of the image stretch used).

6.3 Hot Pixels

Hot pixels represent the largest excursions from the median per-camera dark current rate, much larger in amplitude than the low-order tree ring patterned variations. Table 6 gives statistics about the fraction of pixels in each camera with measured dark current rates exceeding 10×, 100× and 1000×. While hot pixels clearly represent a very small fill factor of the pixels within each array, there are $>$ 6 million total pixels per detector and hence the dark current correction and masking of hot pixels turns out to yield the largest cosmetic image quality improvement of any processing step during CI image reductions.

[10] /project/projectdirs/desi/users/ameisner/CI/post_install_calibs/CI_master_dark.fits

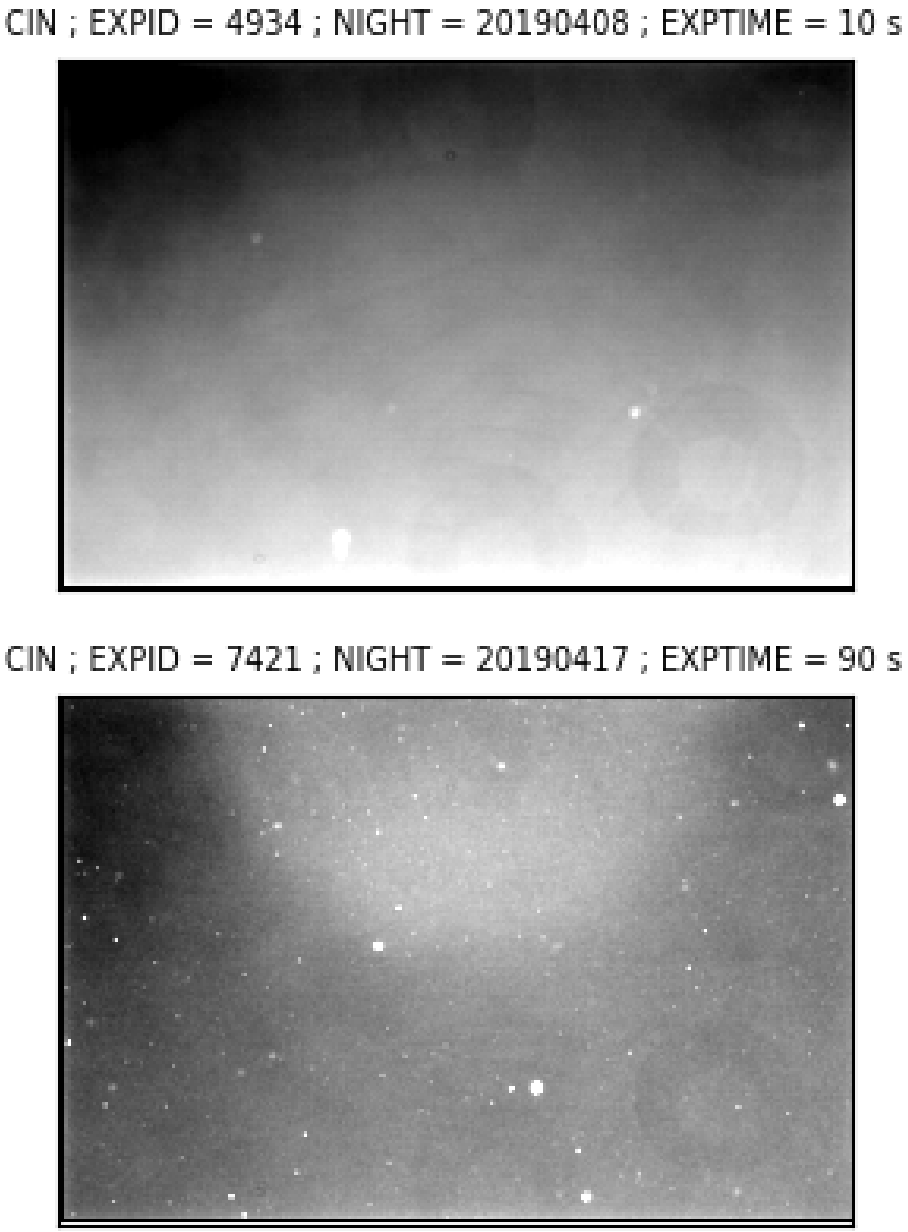

Figure 8: Seemingly "good" twilight flat (top), and twilight flat apparently contaminated by light leaks (bottom). Both panels show CIN, plotted with the same orientation in which upward on the page is radially outward with respect to the focal plane center. The image at top shows primarily a gradient with CI pixel Y coordinate, as would be expected from vignetting, whereas the low-order structure in the image at bottom is dominated by a qualitatively different component. In this colormap, white indicates relatively high signal and black indicates relatively low signal.

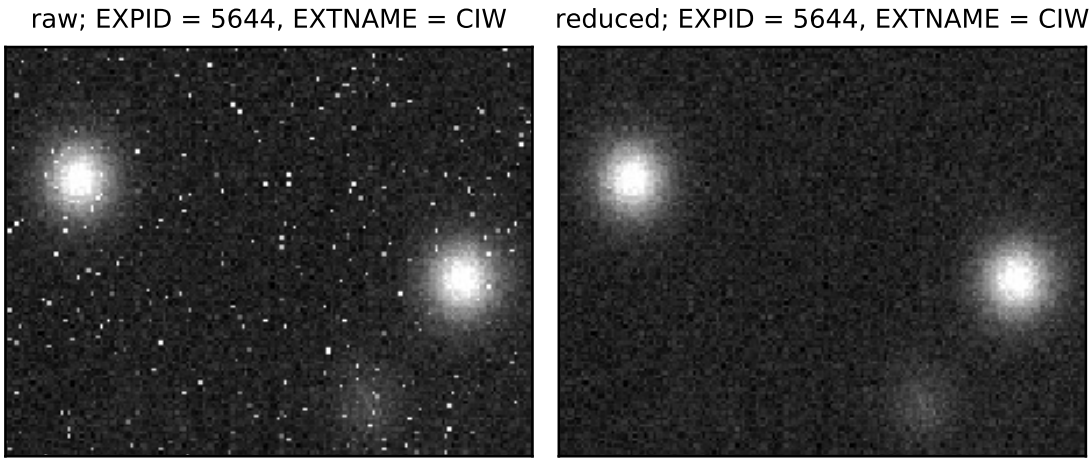

Figure 9: The vast majority of the reduced CI data's improved cosmetic appearance comes from accounting for dark current. This is illustrated by the above comparison of CI image cutouts for a long (1200 second) exposure (EXPID = 5644). The raw data is shown at left, whereas the ci_reduce reduced data is shown at right. Each cutout is $21.4'$ $\times$ $16.6'$ in angular extent. The bright pixels at left not due to astrophysical sources are those with relatively high dark current rates.

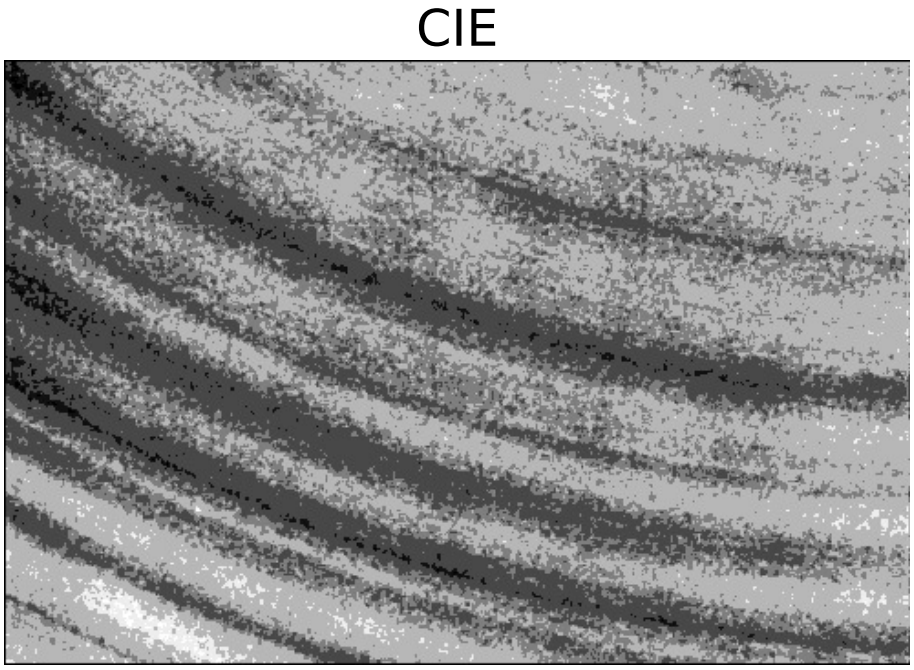

Figure 10: 16 pixel median filtered CIE master dark current image. The grayscale rendering shows the ratio of the median filtered image to its overall median value, with the grayscale ranging from 0.97 (black) to 1.03 (white). Aside from hot pixels, low-level tree rings are the dominant structures seen in the darks.

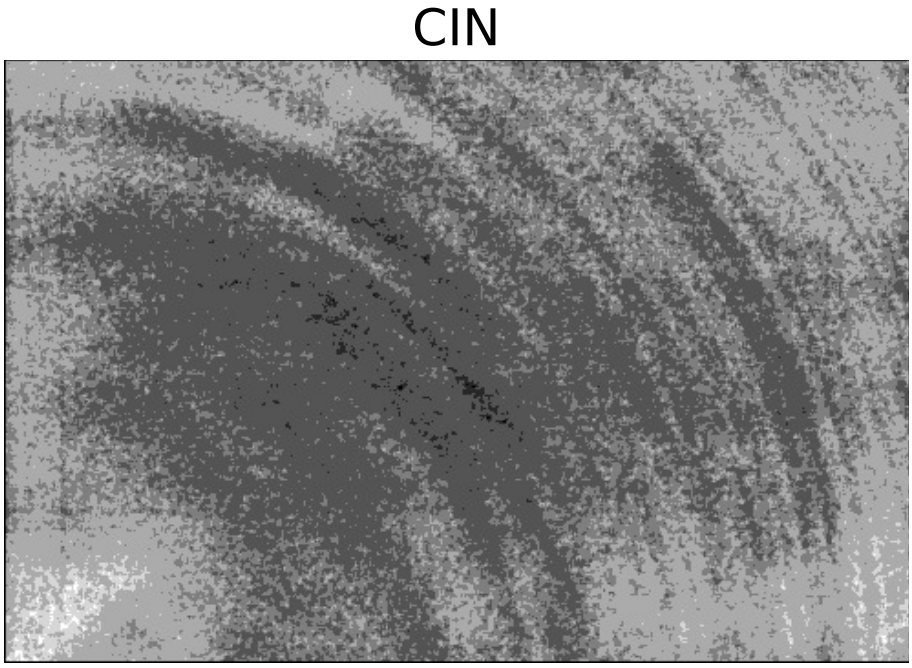

Figure 11: 16 pixel median filtered CIN master dark current image. The grayscale rendering shows the ratio of the median filtered image to its overall median value, with the grayscale ranging from 0.97 (black) to 1.03 (white). Aside from hot pixels, low-level tree rings are the dominant structures seen in the darks.

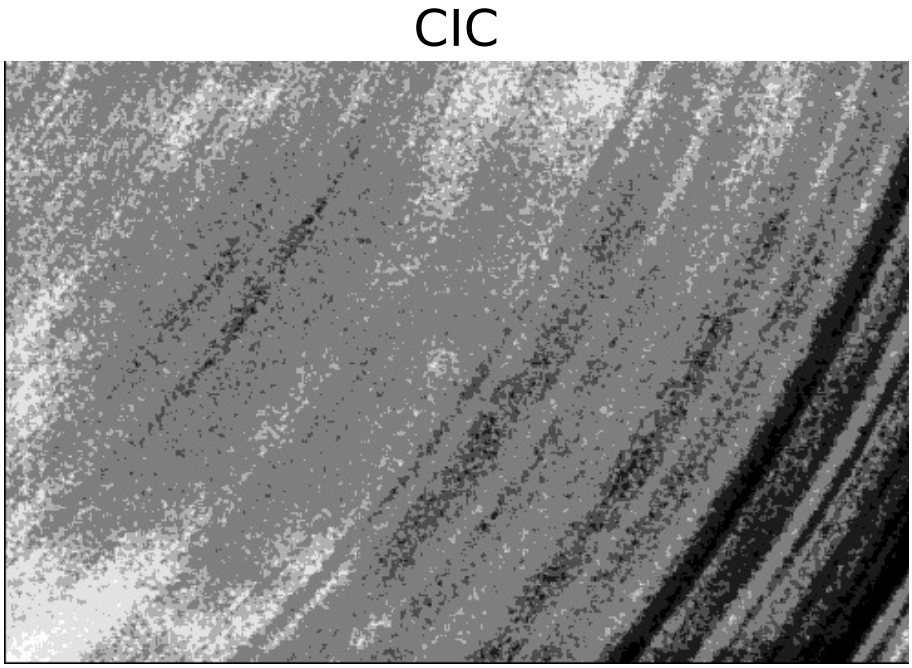

Figure 12: 16 pixel median filtered CIC master dark current image. The grayscale rendering shows the ratio of the median filtered image to its overall median value, with the grayscale ranging from 0.97 (black) to 1.03 (white). Aside from hot pixels, low-level tree rings are the dominant structures seen in the darks.

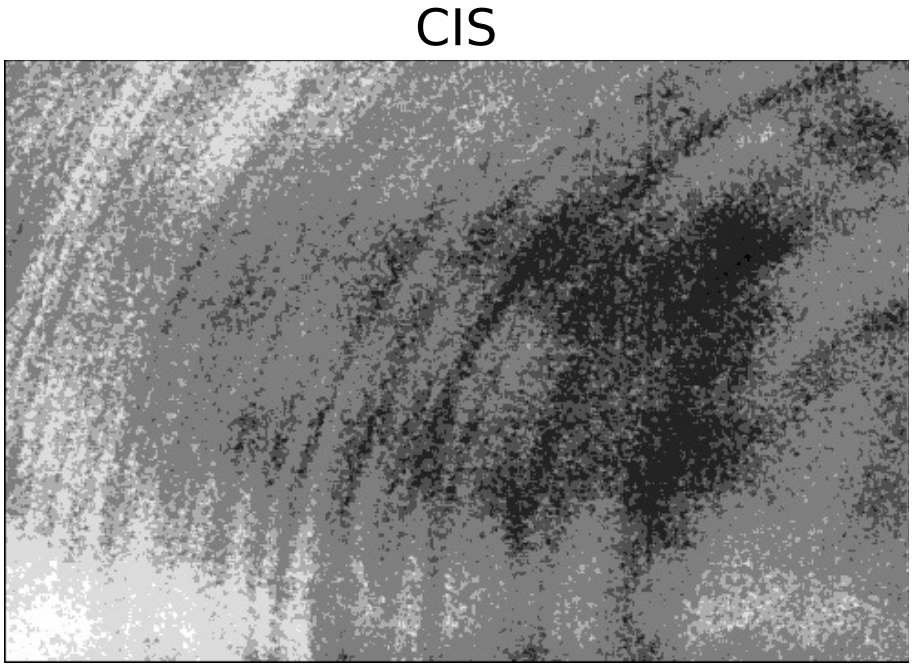

Figure 13: 16 pixel median filtered CIS master dark current image. The grayscale rendering shows the ratio of the median filtered image to its overall median value, with the grayscale ranging from 0.97 (black) to 1.03 (white). Aside from hot pixels, low-level tree rings are the dominant structures seen in the darks.

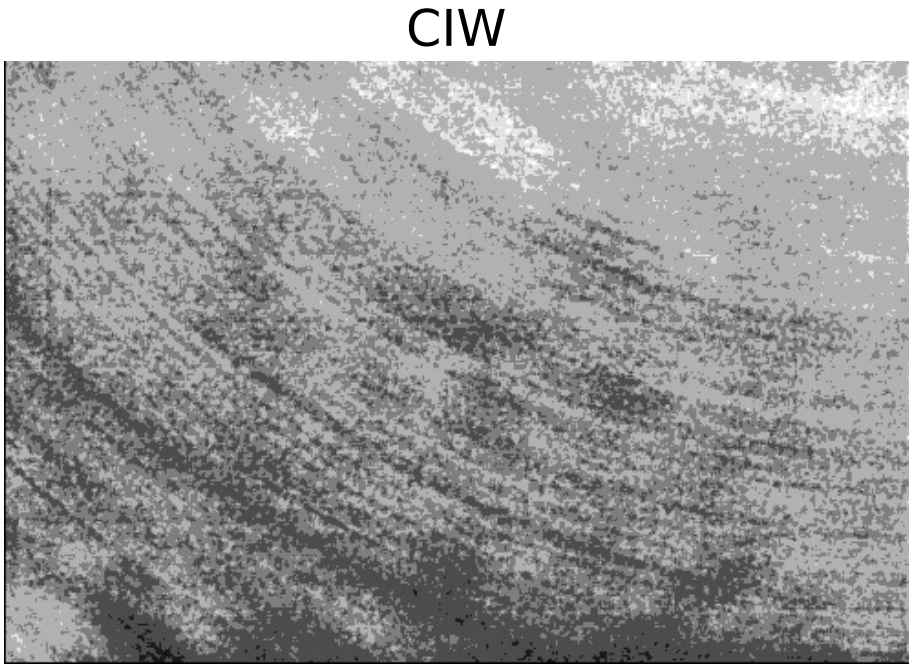

Figure 14: 16 pixel median filtered CIW master dark current image. The grayscale rendering shows the ratio of the median filtered image to its overall median value, with the grayscale ranging from 0.97 (black) to 1.03 (white). Aside from hot pixels, low-level tree rings are the dominant structures seen in the darks.

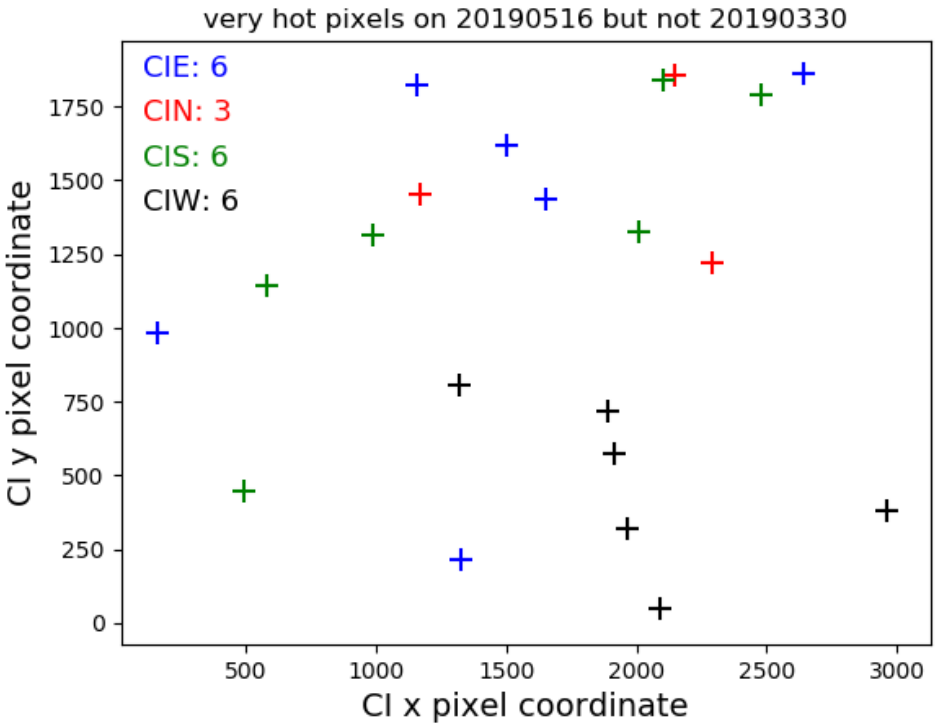

Figure 15: A small number of CI pixels (of order half a dozen per camera) became 'very hot' by mid-May, despite having shown normal levels of dark current at the beginning of the CI run. The locations of the pixels that changed from not hot to very hot over time are labeled here, showing that they are fairly randomly distributed within each camera. Each plus mark represents a single pixel. CIC is not shown but presumably has a few such pixels as well. These pixels are problematic if not corrected or masked because their high dark current levels can lead them to be detected as spurious $r \sim 19$ stars at the exact same pixel location in every exposure.

camera	median (ADU/pix/s)	median ($e-$/pix/s)	assumed gain ($e-$/ADU)	T (° C)	frac >10× median	frac >100× median	frac >1000× median
CIW	0.143	0.251	1.750	7.5	2.86×10^{-3}	1.56×10^{-4}	1.03×10^{-5}
CIC	0.143	0.245	1.710	7.5	2.78×10^{-3}	1.83×10^{-4}	1.40×10^{-5}
CIN	0.162	0.276	1.705	7.5	2.72×10^{-3}	1.48×10^{-4}	1.40×10^{-5}
CIS	0.153	0.258	1.685	7.5	2.73×10^{-3}	1.50×10^{-4}	1.29×10^{-5}
CIE	0.128	0.223	1.735	7.5	2.86×10^{-3}	2.58×10^{-4}	1.68×10^{-5}

Table 6: Dark current summary statistics based on the master dark images described in §6.1. The assumed gain values are adopted based on those of Table 7.

6.3.1 Time variation of hot pixels

There were definitely ∼6-10 pixels per CI CCD that started out "not hot" at the beginning of the run but later became "hot". Noticed these because they created bogus source extractions at the same pixel locations in every image. Would be amusing to know if there are any cases that went the other direction, from "hot" early in the CI run to "not hot" later on.

7 Linearity

8 Gain

Two independent calibration screen gain measurement data sets were taken on 20190411 and 20190527, near the beginning and end of the CI run respectively. The purpose of this exercise is to check how repeatable the 20190411 gain measurement is using a new data set gathered many weeks later. This will inform our understanding of the systematic uncertainty on throughput measurements/comparisons.

I made a rough initial pass at measuring the CI camera gains from the calibration screen data that I took on 20190411. Here are the results that I get:

camera	20190411	20190527	adopted	A. Ross [desi-commiss 602]
CIE	1.73	1.74	1.735	1.67
CIN	1.70	1.71	1.705	1.64
CIC	1.71	1.71	1.71	1.65
CIS	1.68	1.69	1.685	1.61
CIW	1.75	1.75	1.75	1.64

Table 7: Summary of gain measurements. The "20190411" and "20190527" columns are those derived in this work using Mayall/DESI calibration screen data. The values in the rightmost column are based on pre-installation CI engineering data and assume that the reported gain measurements from [desi-commiss 602] had camera labels swapped in the same way as did the 20190401 on-sky CI data.

Summary plots are attached. This is only a rough first cut, but the values are indeed fairly close to the 1.64 e-/ADU value that I assumed previously in my throughput/zeropoint calculation from [desi-commiss 850]. I ignored details of the bias and readnoise by performing linear fits with the y intercept left as a free parameter, since the readnoise is just an offset in the y axis and the bias is just an offset in the x axis. In the future I'll try actually accounting for these in detail. I'll be interested to see whether these or other detailed fitting/procedural choices can matter at the percent level. I'm encouraged by how linear the trend of variance with counts is over the range I've plotted. Regarding, the 1.2 value from [desi-commiss 959], that's different from what I get by a suspicious factor of sqrt(2). If it does turn out that the gains are slightly higher than previously assumed, (1.7 e-/ADU instead of 1.65 e-/ADU) then the throughput without aperture mask was actually a few percent better than what I initially reported.

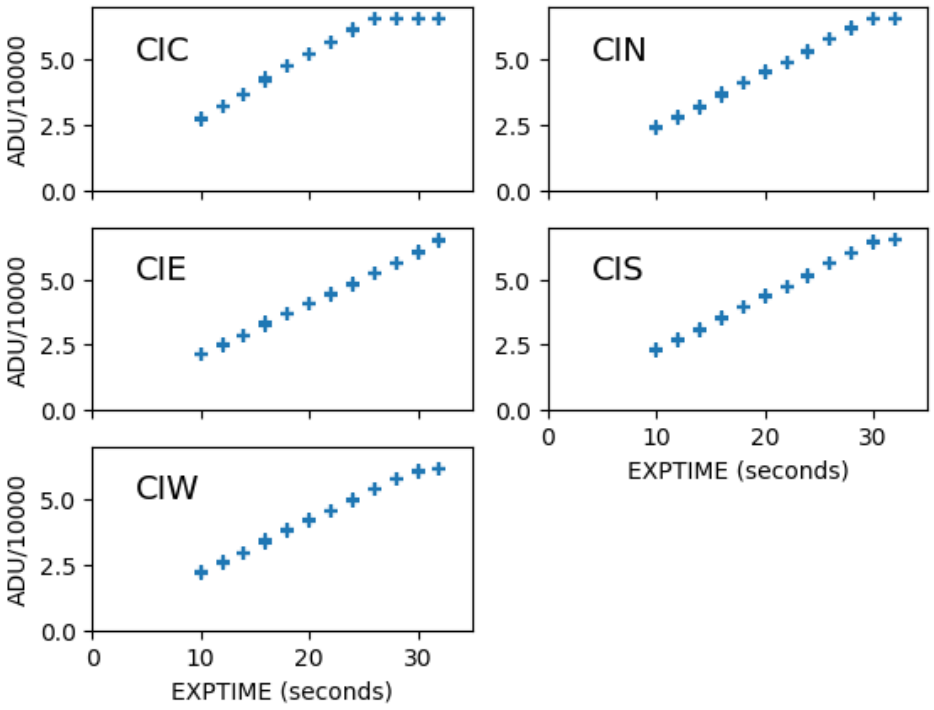

Figure 16: Median counts versus exposure time from the same calibration screen data used to measure the gain. The low dome lights are not intended to provide an illumination level that is perfectly constant in time, yet these plots remain consistent with linear response up to roughly the nominal full well (100,000 $e-$) quoted by the CI CCD data sheet.

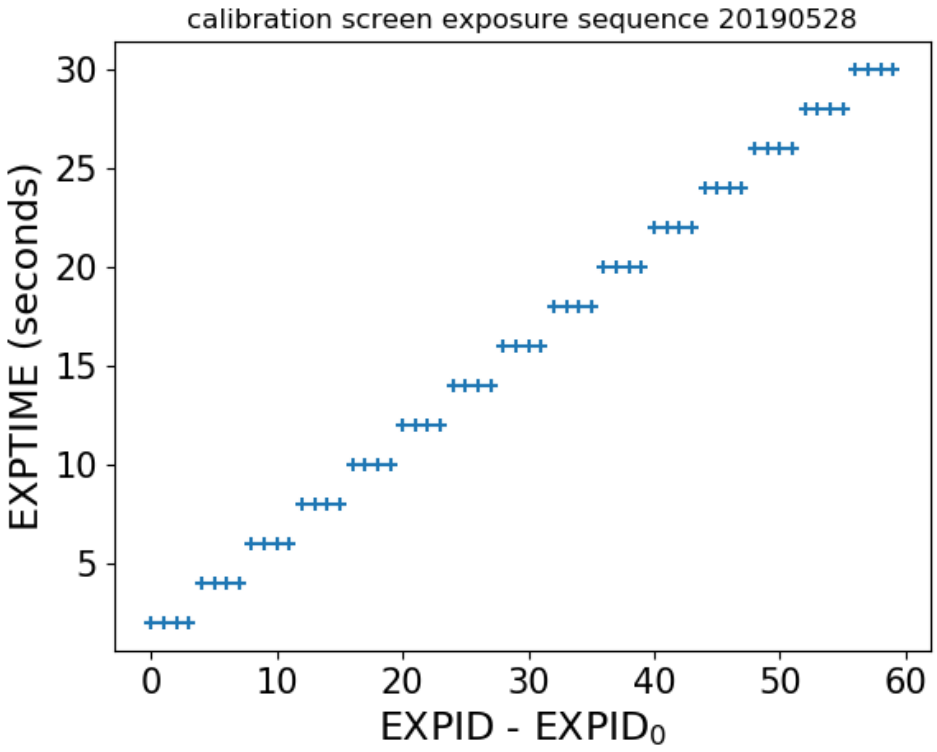

Figure 17: Summary of the exposure sequence used to measure the gain via images of the Mayall/DESI calibration screen with the low dome lights on. At each unique exposure time value, a set of 4 consecutive exposures are obtained, allowing for six pairs of difference images with which to measure the variance at each count level.

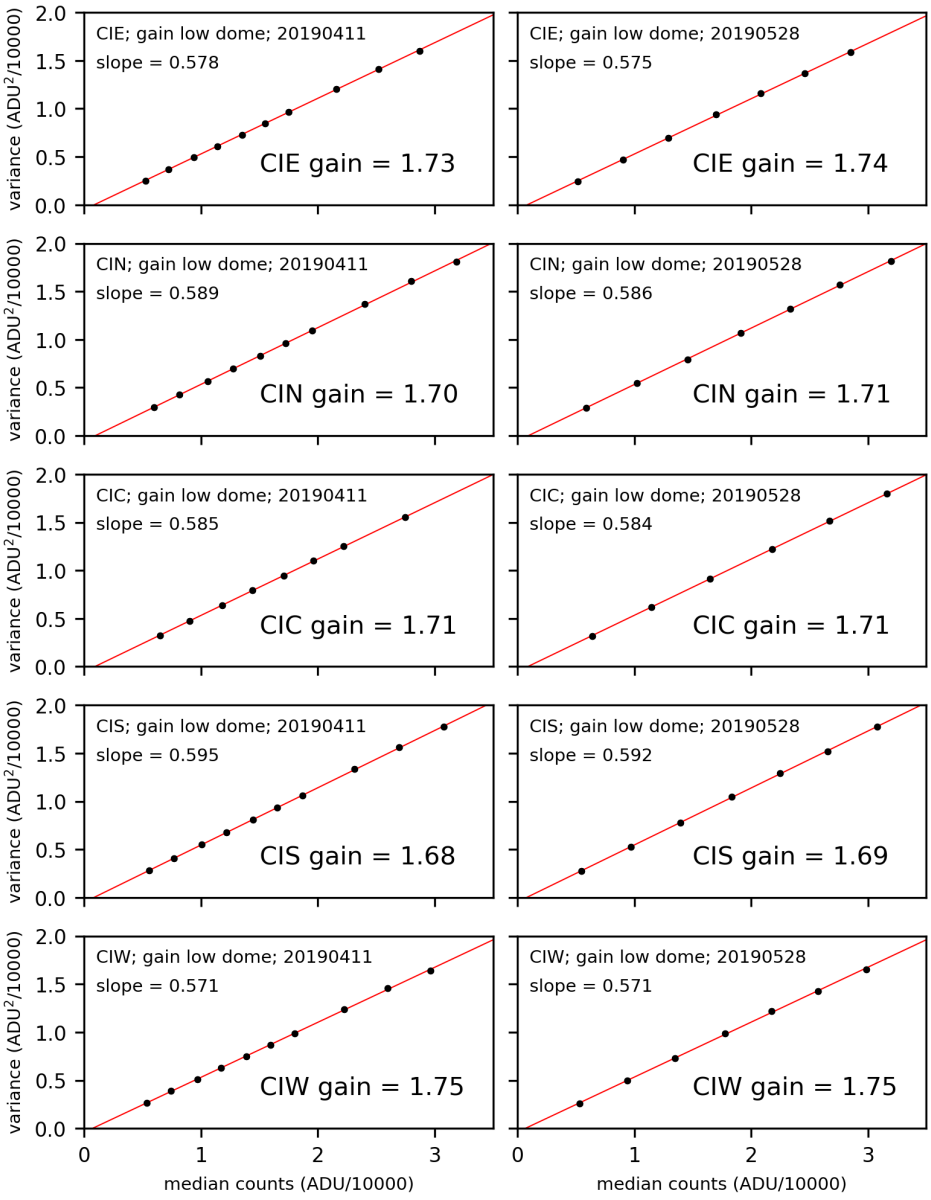

Figure 18: Summary of gain measurements.