



Vera C. Rubin Observatory
Systems Engineering

LSST Camera Electro-Optical Test (Run 7) Results

LSST Camera Team, Pierre Antilogus, Pierre Astier, John Banovetz,
Johan Bregeon, James Chiang, Seth W. Digel, Kevin Fanning,
Yassine Faris, Thibault Guillemain, Anthony S. Johnson,
Claire Juramy-Gilles, Craig S. Lage, Shuang Liang,
Sean Patrick MacBride, Stuart Marshall, Homer Neal, Daniel Polin,
Andrew P. Rasmussen, Aaron Roodman, Eli S. Rykoff, Brian Stalder,
John Gregg Thayer, J. Anthony Tyson, Hye Yun Park,
Yousuke Utsumi, and Zhuoqi Zhang

SITCOMTN-148

<https://doi.org/10.71929/rubin/2583999>

Latest Revision: 2025-08-19

DRAFT

Abstract

This note collects results from the LSST Camera electro-optical testing prior to installation on the TMA. We describe the focal plane characterization, optimization, and the resulting default settings. Results from eo-pipe are shown for standard runs such as B-protocols, Dense and SuperDense PTCs, gain stability, OpSim runs of Darks, and Darks with variable delays. We also describe features such as e2v Persistence, ITL phosphorescence in coffee stains, remnant charge near Serial register following saturated images, vampire pixels, ITL dips, and other sensor features.

Draft

Change Record

Version	Date	Description	Owner name
1	YYYY-MM-DD	Unreleased.	Yousuke Utsumi

Document source location: <https://github.com/lsst-sitcom/sitcomtn-148>

Draft

Contents

1 Executive Summary	1
2 Introduction	6
3 Electro-optical setup	9
3.1 Run 7 optical modifications	9
3.2 Projector spots	10
3.3 Dark current and light leaks	14
3.3.1 Light leak mitigation with shrouding the camera body	14
3.3.2 Filter Exchange System Autochanger light leak masking	15
3.4 Shutter condition impact on darks	17
3.5 Impact of filters on dark images	18
4 Reverification	21
4.1 Background	21
4.2 Stability flat metrics	22
4.2.1 Charge transfer inefficiency	22
4.3 Dark metrics	27
4.3.1 Dark current	27
4.3.2 Bright defects	29
4.4 Flat pair metrics	31
4.4.1 Linearity and PTC turnoff	31
4.4.2 PTC gain	35
4.4.3 Read Noise	36
4.4.4 PTC Noise	36
4.4.5 Brighter-fatter coefficients	37
4.4.6 Row-means variance	43
4.4.7 Divisadero tearing	45
4.4.8 Dark defects	47
4.5 Persistence	50

4.6	Differences between Run 6 and Run 7	53
5	Camera Optimization	54
5.1	Persistence optimization	54
5.1.1	Persistence optimization	56
5.1.2	Impact on full-well	57
5.1.3	Impact on brighter-fatter effect	58
5.2	Sequencer optimization	59
5.2.1	Improved clear	60
5.2.2	Not toggling the RG bit during parallel transfer for e2v sensors	66
5.2.3	Disabling IDLE FLUSH	69
5.3	Summary	71
6	Characterization & Performance Stability	72
6.1	Final characterization	72
6.1.1	Background	72
6.1.2	Stability flat metrics	72
6.1.3	Dark metrics	76
6.1.4	Flat pair metrics	78
6.1.5	Persistence	95
6.1.6	Differences between Run 7 initial and Run 7 final measurements	97
6.2	List of problematic amplifiers	99
6.3	Defect stability	102
6.3.1	Bright defects	104
6.3.2	Dark defects	104
6.4	Bias stability	105
6.5	Glow search	111
6.6	Illumination-corrected flat	113
6.7	Full well measurements	116
6.7.1	Spot photometry model	117
6.8	Non-linearity studies	120

6.9	Gain stability	122
6.9.1	Non uniformity in focal plane response change	129
6.9.2	Non uniform temperature variation.	129
6.9.3	ITL gain glitches	133
6.10	Guider operation	134
6.10.1	Guide mode	134
6.10.2	Guider timing	135
6.10.3	Noise investigation	136
6.10.4	Impact on science sensors	137
6.11	Summary	140
7	Sensor Features	142
7.1	Tree rings	142
7.1.1	Centers of the tree rings	142
7.1.2	Radial study	142
7.1.3	Effect of diffuser	143
7.1.4	Voltage dependency	144
7.1.5	Wavelength dependence	145
7.2	ITL dips	149
7.3	Vampire pixels	154
7.3.1	First observations	154
7.3.2	LSST Camera vampire pixel features	155
7.4	Phosphorescence	159
7.4.1	Measurement techniques for detecting and quantifying phosphorescence	160
7.4.2	Results of phosphorescence detection in ITL sensors	162
7.4.3	Summary of phosphorescence in ITL sensors	169
7.5	Summary	172
8	Overall Operation and Issues	174
8.1	Timeline of events and issues	174
8.2	Camera control network performance	176

8.3	REB PS power trip	178
8.4	FES latch sensor failure during motion	179
8.5	PCS degradation	180
8.6	R24/Reb0 and UT leak fault issue	182
8.7	Data corruption	183
8.8	Guider high gain issue	184
9	Conclusions	185
9.1	Run 7 final operating parameters	185
9.1.1	Voltage conditions	185
9.1.2	Sequencer conditions	185
9.2	Record runs	186
9.3	Other runs of relevance	187
A	Reference figures	195
A.1	From Sections 4 and 6.1	195
A.1.1	Focal-plane measurement figures	195
A.1.2	Differential histograms	196
A.2	Web report reference figures	197
A.2.1	Focal plane level	197
A.2.2	Raft level	200
A.2.3	Sensor level	205
B	Observatory Control System integration	213
B.1	LSSTCam in OCS	213
B.1.1	Configurable SAL Components: base unit of the OCS	213
B.1.2	ts_xml	215
B.1.3	OCS Scheduler and Script Queue	215
B.1.4	OCS-Bridge	216
B.2	Testing steps and achievements in Run 7	217
B.3	Further testing and development	218

C	Phosphorescence	219
D	Phosphorescence identification on a set of ITL sensors	221
E	Phosphorescence morphological comparisons with features seen in <i>blue</i> flat field response	233
F	Phosphorescence kinetics characterization	240
G	Phosphorescence response characterization	249
H	References	259
I	Acronyms	261

LSST Camera Electro-Optical Test (Run 7) Results

1 Executive Summary

Introduction

The LSST Camera (LSSTCam; SLAC National Accelerator Laboratory & NSF-DOE Vera C. Rubin Observatory, 2025) was constructed at the SLAC National Accelerator Laboratory in Menlo Park, California, USA. Its functionality and performance were evaluated through various integration phases, leading to the identification and mitigation of non-ideal features. In May 2024, LSSTCam was transported from Menlo Park to Cerro Pachón, Chile, where the Vera C. Rubin Observatory is being constructed. LSSTCam was installed in the clean room (White Room) on Level 3 of the observatory. After the power and cooling lines were connected and the vacuum performance verified, the LSSTCam underwent the seventh series of electro-optical (EO) testing (Run 7) from September 2024 to December 2024, collecting 56,066 exposures. This report describes the results from Run 7.

Key Testing Points

- *Testing Setup Differences:* The EO test setup in the White Room differed from previous setups, primarily in using the CCOB Wide Beam projector with a new diffuser system.
- *Performance Verification Post-Transportation:* The Camera's performance was reverified after transportation to ensure it matched the pre-transportation checks.
- *Optimization of Features:* Previous EO testing identified issues such as persistence and bias instability, which were optimized during Run 7.
- *Camera Performance Post-Optimization:* The Camera's performance was evaluated after implementing optimizations.
- *Investigation of Other Features:* Additional sensor features were investigated.
- *Summary of Run 7 Operations and Issues:* Overall operations and issues encountered during Run 7 were summarized.

Electro-Optical Setup

- *Run 7 Optical Modifications:* The EO test setup in the White Room included a cone attached to the L1 cover and a shroud to create a dark environment. A diffuser was installed to reduce weather patterns from turbulence of the air between the L1+L2 lenses and L3 and uniformly illuminate the focal plane.
- *Projector Spots:* A 4K projector was used for EO testing, illuminating all 3206 amplifiers with individual spots. The projector's background illumination posed challenges, but adjustments were made to improve contrast.
- *Dark Current and Light Leaks:* Initial measurements identified light leaks, which were mitigated by covering gaps and using a blackout fabric shroud.

Reverification

- *Background:* EO camera test data was processed to extract key metrics. The primary concern was maintaining performance characteristics between Run 6 and Run 7.
- *Stability Flat Metrics:* Charge transfer inefficiency (CTI) measurements showed consistent performance between runs.
- *Dark Metrics:* Dark current measurements indicated improved performance, likely due to better shrouding.
- *Bright Defects:* Bright defects were evaluated, showing a small increase in Run 7.
- *Flat Pair Metrics:* Linearity and PTC turnoff metrics were consistent, with minor differences for e2v sensors due to voltage changes.

Camera Optimization

- *Persistence Optimization:* Persistence issues were addressed by adjusting operating voltages, significantly reducing residual signals.

- *Impact on Full-Well:* Reducing the parallel swing voltage decreased the full-well capacity by 23%.
- *Impact on Brighter-Fatter Effect:* The brighter-fatter effect increased slightly but remained within acceptable levels.
- *Sequencer Optimization:* Several sequencer configurations were tested to improve performance, including changes to the clear method and toggling the RG bit during parallel transfer.

Characterization & Performance Stability

- *Final Characterization:* Key metrics from initial and final Run 7 configurations were compared, showing high consistency.
- *Stability Flat Metrics:* Serial and parallel CTI measurements remained stable.
- *Dark Metrics:* Dark current and bright defect measurements were consistent, with improvements in some rafts.
- *Flat Pair Metrics:* Linearity and PTC turnoff metrics showed minor changes, primarily in e2v sensors.
- *PTC Gain:* PTC gain measurements were consistent, with a slight increase in e2v sensors.
- *Read Noise:* Read noise remained stable across runs.
- *PTC Noise:* PTC noise measurements showed no significant deviations.
- *Brighter-Fatter Coefficients:* The brighter-fatter effect increased slightly in e2v sensors.
- *Row-Means Variance:* Row-means variance showed a slight decrease in e2v sensors.
- *Divisadero Tearing:* Divisadero tearing was significantly reduced in e2v sensors.
- *Dark Defects:* Dark defect counts remained consistent.
- *Persistence:* Persistence was minimized in e2v sensors, with sub-ADU levels across the LSST bandpass.

Sensor Features

- *Tree Rings*: Tree rings are concentric variations in silicon doping concentration observed in flat images. The centers of the tree rings were measured for all 189 LSST Camera science sensors, showing consistent positions around the average center for each direction.
- *ITL Dips*: ITL dips were investigated using spot and rectangle projections. No evidence of ITL dips was found in the lab data, but further investigation is needed for on-sky data.
- *Vampire Pixels*: Vampire pixels are a group of pixels with high photo-response surrounded by pixels with low photo-response. These features were identified on ITL sensors and correlated with phosphorescence.
- *Phosphorescence*: Phosphorescence was observed in some ITL sensors, showing a transient signal after exposure. The effect was dependent on the HV bias state and varied with wavelength and signal level.

Operations and issues

- *Camera Control Network Performance*: Network issues were addressed by simplifying the network configuration and upgrading the control system.
- *REB PS Power Trip*: Power trips were mitigated by grounding the Utility Trunk door and implementing ESD controls.
- *FES Latch Sensor Failure*: A faulty cable was replaced, resolving the issue.
- *PCS Degradation*: Performance degradation was managed through various mitigation strategies, with further analysis required.
- *R24/Reb0 and UT Leak Fault*: Low temperature issues were resolved by adjusting the Dynalene temperature and adding a load resistor to the DC-DC converter of the leak detector.
- *Data Corruption*: Data corruption issues were resolved by restarting Data store RCEs.

- *Guider High Gain Issue:* The guider issue was resolved by power cycling and resetting the RCE.

Run 7 successfully demonstrated LSSTCam's readiness for installation on the Telescope Mount Assembly (TMA). The optimizations and mitigations are implemented balancing both the required specifications and the real scientific needs, with robust performance across various metrics. The tests highlighted a few potential issues that need to be addressed.

Draft

2 Introduction

The LSST Camera (LSSTCam) was constructed at the SLAC National Accelerator Laboratory in Menlo Park, California, USA (Lange et al., 2024). The functionality and performance of the Camera had been studied in various integration phases from two rafts testing (Run 1; March 2019–April 2019 and Run 2; June 2019), nine rafts testing (Run 3; Oct 2019–Nov 2019), the full focal plane testing (Run4; Aug 2020–Nov 2020 and Jan 2020–Feb 2021), the full focal plane with the utility trunk (Run 5; Nov 2021–Jan 2022), and the full Camera testing (Run 6a; June 2023 and Run 6b; Oct 2023–Oct 2024)(Roodman et al., 2024). These testing periods verified the Camera functionality and led to discoveries of non-ideal features and means of mitigating many of them (Utsumi et al., 2024).

In May 2024, LSSTCam was loaded into a Boeing 747 airplane in San Francisco, CA, flown by air, and then transported by trucks from Santiago, Chile to the summit of Cerro Pachón at 2700 m in the Andes mountain range in Chile, where the NFS-DOE Vera C. Rubin Observatory is being constructed. The LSST Camera was transferred from the truck to a support structure which was then rolled into the clean room, known as the White Room, on Level 3 of the Vera C. Rubin Observatory. After connecting lines for power and cooling and verifying the integrity of the vacuum, cool down began in late August. Following this, the seventh series of electro-optical (EO) testing, Run 7, prior to installation on the Telescope Mount Assembly (TMA), was conducted from the end of September 2024 to the beginning of December 2024 to reverify its performance and undertake further optimization. We collected 56,066 exposures during this testing campaign since the Camera became the full operational state (high voltages applied to CCDs) and the data were sent to the processing nodes at the SLAC Shared Scientific Data Facility (**S3DF**).

This document details initial interim testing results, with a focus on several points:

- What is the difference of testing setup? (Section 3)
- Does the Camera after the transportation still perform as we checked out in California? (Section 4)
- Optimizations to the features that we found during previous EO testings such as persistence and bias instability. (Section 5)
- How does the Camera perform after implementing those optimizations? (Section 6)

- Investigating other features (Section 7)
- Summarize the overall operations and issues during Run 7 (Section 8)

All the results presented here are subject to future changes.

Figure 1 shows the layout of the focal plane. The LSST focal plane consists of 21 Science Rafts and 4 Corner Rafts (cyan). Science Rafts have two varieties based on the vendor of the sensors used for the Raft: ITL (green) or e2v (yellow). Each sensor has 4kx4k pixels, segmented by 16 channels to make the fastest readout being 2 seconds. Corner Rafts have two different kinds of sensors: guider sensors and a wavefront sensor. Guider sensors are ITL sensors as are some of the other science sensors, while the wavefront sensor is an ITL sensor split in the middle and packaged into one sensor with an offset by $\pm 2\text{mm}$ with respect to the other science sensors to provide off-focus point source images to measure “donuts” – an image of the pupil used to regulate the focus of the telescope.




Figure 1: The focal plane layout

3 Electro-optical setup

3.1 Run 7 optical modifications

For Run 7 in the White Room on Level 3 our EO test setup had a few differences from the Run 6 setup in IR2 at SLAC. One difference was that we were not able to use the CCOB Narrow/Thin beam because we did not have the resources to configure it. As such, the majority of the testing was done with the CCOB Wide Beam projector (Utsumi et al., 2024). We did obtain an additional projector, the 4K projector, partway through Run 7 that will be discussed later. With the CCOB Wide Beam, we used a cone attached to the L1 cover as well as a shroud to create a dark environment (Fig. 2).




Figure 2: (left) Final shroud configuration of LSSTCam in Level 3 to reduce light leaks. (right) CCOB Wide Beam attached to the cone and shrouded.

This allowed us to operate on Level 3 with a dark current of <0.1 ADU/sec with the shutter open. The initial setup of the CCOB Wide Beam projector was the same as for Run 6, with a minimal ND filter (10%) attached to a C-mount lens. One difference was that the f/stop of the lens was changed from 2.6 to 1.6 (fully open). This was done to try to reduce the effect of the ‘weather’ and the ‘CMB pattern’ two effects that we found in Run 6 and were found to be due to our projection setup (see Banovetz et al. (2024)). While changing the f/stop did reduce the weather pattern, it also caused a much steeper illumination roll-off across the focal plane (see Figures 5 and 6). We evaluated the weather pattern and illumination roll-off relative to Run 6.

To both reduce the effects of the ‘weather’ and ‘CMB’ but retain uniform illumination across the focal plane, we installed a diffuser in the cone attached to L1. Figure 3 shows the placement of the diffuser within the cone. The diffuser used is a 60° diffusing angle unmounted sheet from Edmunds Optics.




Figure 3: Diffuser installed into the light cone.

We found that the diffuser greatly reduced the ‘weather’ (Fig. 4) and eliminated the CMB pattern and more uniformly illuminated the focal plane (Fig. 5), and removed the extreme roll-off (Fig. 5 and Fig. 6) with a penalty of decreasing the overall illumination by roughly 35% even though we fully opened the f-stop.

The diffuser was installed for all B protocol and PTC runs (see Sec. 4) moving forward, being taken out only for pinhole projection runs and when using the 4K projector.

3.2 Projector spots

The addition to the projectors used for EO testing was a 4K projector (Epson LS11000 LCD), similar to those used in conference rooms. This projector was first tested at SLAC and arrived

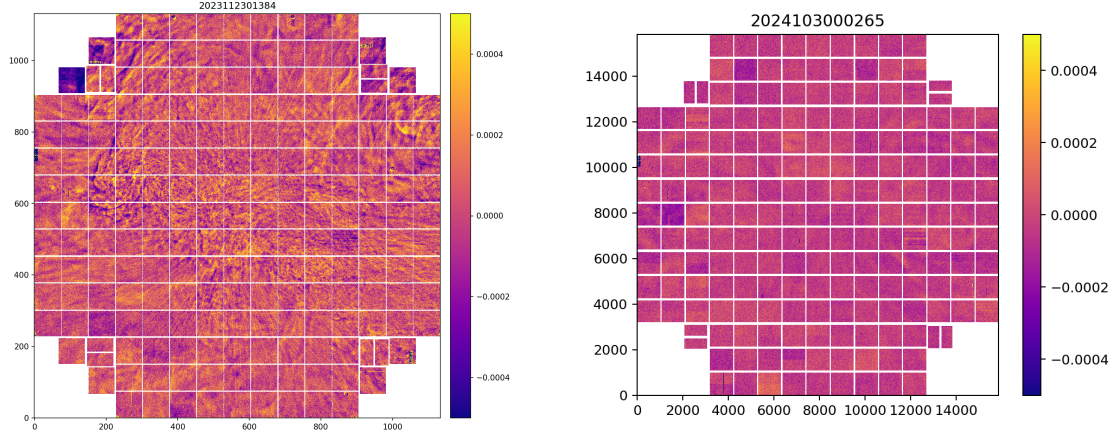


Figure 4: Full focal plane fractional difference of images for Run 6 (left) and Run 7 (right).




Figure 5: (left) Illumination across the focal plane from Run 7 without the diffuser (E968) as compared to Run 6. (right) Illumination across the focal plane from Run 7 with the diffuser (E1047) as compared to Run 6.




Figure 6: Combined flat images from three different runs that show the roll-off of the CCOB Wide Beam with Run 7 data without the diffuser and with the projector aperture open (left), Run 7 data with the diffuser and projector aperture closed (middle), and Run 6 without the diffuser but the projector aperture closed (right). The pixels are normalized and the left figure is missing detectors as we were still in the process turning them off when doing this test.

at the observatory about halfway through Run 7. It was used primarily as a spot projector, as the pinhole filter was not available at that time because of the Filter Exchange System was temporarily inoperable. The projector has an advantage, instead, as it could illuminate all 3206 amplifiers instead of just the 21 illuminated by the pinhole projector. Figure 7 shows both the setup of the projector on Level 3 and an example of a spot image and the spots across the focal plane. Since the projector does not have fast illumination control, we primarily used the LSSTCam main shutter instead of any flashing of the light source (e.g., as we did with the LEDs of the CCOB Wide Beam). One downside that was found was that the projector illuminated the entire focal plane at some background level, not just the spot regions. The background illumination also had structure that changed with time and could not be easily subtracted. Figure 8 shows an example of a spot image of just one detector as well as a zoomed in image of a single spot which highlights the background structure. The resulting contrast between the spot and the background was only about a factor of 6. Changing the spot shape to large rectangles for crosstalk measurements increased the contrast ratio to 30. Examples of the rectangles can be seen in Figure 9. Though the contrast was much improved, there was still a background structure as can be seen in the saturated image of the figure.

This section describes the spots and rectangle patterns used for tests with the 4K projector.

- Projector background
- Spots on many amps

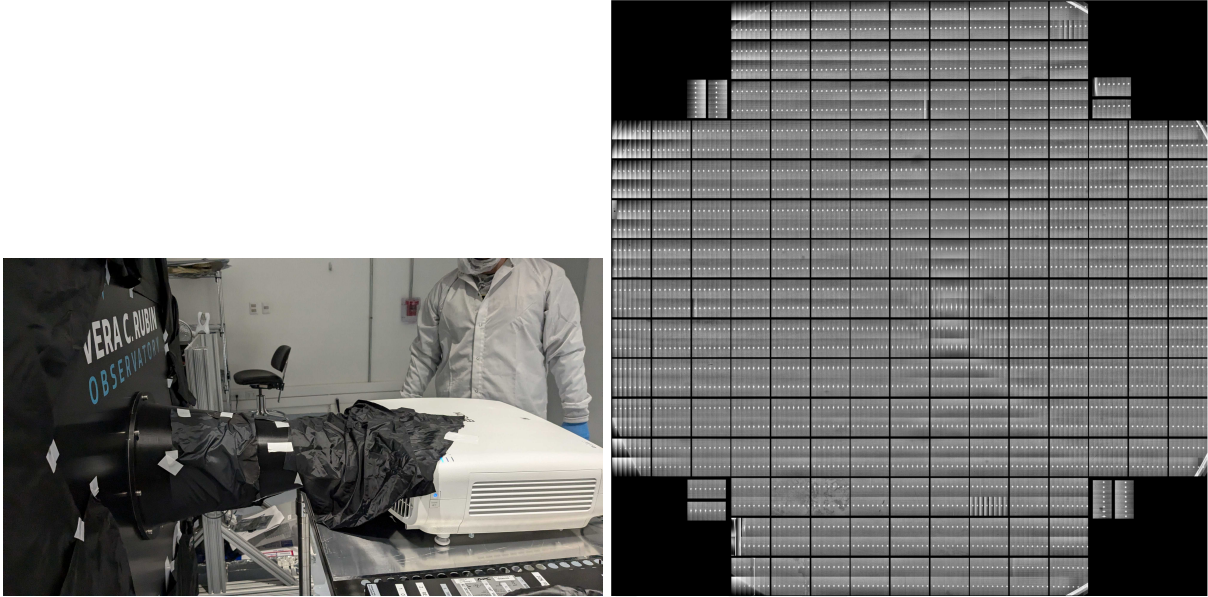


Figure 7: (left) The spot projector set up on Level 3. (right) An example of an image taken with the spot projector with all the amplifiers containing a spot.




Figure 8: (left) Example of a spot image zooming into a single detector. (right) Example of a spot image zooming further into a single spot. In both the images, there is a clear background structure caused by the projector.




Figure 9: (left) Example of a spot image that utilized the rectangle shape, zoomed into a single detector (left), zoomed into the spot (middle) and zoomed into the spot with a saturated image to highlight the background pattern caused by the projector (right).

- Spots on one amp
- Optical setup

3.3 Dark current and light leaks

3.3.1 Light leak mitigation with shrouding the camera body

One of the first tests we attempted with LSSTCam was measuring dark current and sources of light leaks in the camera body. Before beginning we covered gaps between the L1 cover and the gaskets with tape, in accessible locations . Figure 10 shows the gaps that we could see between L1 and its cover. The inaccessible locations were later covered with shroud.

Once these were sealed, we took some initial measurements and then covered the LSSTCam body with a Thorlabs blackout fabric shroud (BK5). Figure 2 shows the final configuration of the shroud covering the camera. We also found light leaks where the light cone attached to L1 was housed, and from the Utility Trunk, which were covered with shroud. Table 1 includes the observations, the corresponding measured dark currents, and comments on what changed during the chasing of the leaks.

Table 1: Summary of the 15 s dark exposures, the different conditions, and the resulting dark current. Exposure ID is preceded by “MC_C202409”. The shroud was in place for each of these measurements. (“Initial Covering” was just the CCOB cone and around the L1 cover.)

Exposure	Dark Current (ADU/s)	Room Lights	Shutter	Comments
09_000012	0.16	Off	Closed	
09_000018	0.16	On	Closed	

09_000038	2.94	On	Open	Initial Covering
09_000054	1.34	On	Open	+ Blanket over the FCS
09_000072	0.41	On	Open	+ Blanket over AND under the FCS
09_000078	0.18	Off	Open	+ Blanket over AND under the FCS
10_000031	0.03	On	Open	+ Blanket over AND under the FCS + UT

3.3.2 Filter Exchange System Autochanger light leak masking

A dedicated light leak study of the Filter Exchange System (FES) Autochanger (AC) was performed during Run 6 at SLAC in summer 2023 and a localized faint light source of up to ~ 0.04 $e^-/s/pix$ was found to be associated with the 24 V Clean of the AC.

In the AC this voltage is used to power some probes and all controllers. In February 2024, as AC-1 was extracted from LSSTCam for global maintenance, a direct investigation to localize the light source was performed unsuccessfully. A light source in the AC was not expected, as in the AC all controllers' LEDs have been removed, and most electronics are in "black boxes". Still, two small probes, which had LEDs that could not be removed, were initially masked by a black epoxy. As we had doubts about the quality of this masking at IR wavelengths, we applied extra masking (aluminum black tape) on them during the Feb 2024 maintenance (on AC 1 and 2).

At the start of Run 7 a new study of the light leak based on 900 s dark exposures with the shutter open and the empty frame filter in place, showed that the AC light leaks were still present (see left-hand image of Fig. 12). Following this finding, a full review of all the AC hardware powered by the 24 V dirty was performed, and a candidate was found: the encoders of the five main motors of the AC had only partial documentation from the vendor that did not mention the presence of LEDs. After interaction with the vendor, the encoders were understood to contain ~ 700 nm LEDs. The hypothesis of ~ 700 nm LED sources has been found compatible with the observation as no AC light leaks were detected using various filters (g, r, and y ; all opaque at 700 nm) in LSSTCam at the start of Run 7 (g, r, and y filters). A dedicated test in Paris using an AC spare encoder and a precision photometric set-up allowed identification of the leak in the masking of those LEDs in the vendor packaging (see Figure 11). A complementary masking method based on a 3D printed part + tape + cable tie was qualified in Paris. It was found to mask the light leak and to be safe (all parts correctly secured, see Figure 11)).



Figure 10: Example photos of the L1 cover gaps. These were covered by tape where we could safely apply it and by the black shroud.

In November 2024, we masked all the lights in the back of the Level 3 white room (not the part containing LSSTCam) to set up a high-quality dark room allowing a direct observation with a CMOS camera of the light leak on the AC2 motor encoders. The level of darkness reached allowed us to validate the quality of the light masking of the AC encoders. Notice that the FES-prototype in Paris does not have encoders on the Online Clamps, so we had to tune/qualify the masking of those encoders directly on the AC 2 at the summit.

For both AC 1 and 2, the encoders of the five motors with the vendor issue on their LED masking have been successfully enveloped in a light-tight mask (see Figure 11).

We note that the AC was turned off starting on 27 September 2024 at 21:15 UTC in the first part of Run 7. For the second part of Run 7 (i.e., after mid-November) the AC was back on: as the AC 1 was back in LSSTCam with the new light masks in place on the motor encoders, we were able to take a new series of 900 s darks with the AC turned on and off, confirming that the light leak associated with the FES was eliminated (see right-hand image of Fig. 12).




Figure 11: On the left a picture of one AC encoder. In the center a 60s exposure, with the same encoder turned on, light in the room off: light leak is obvious, mainly associated with holes on the top of the encoder cover. On the right, the X-Online encoder after masking.

3.4 Shutter condition impact on darks

Two runs, E1075 and E1076, were dedicated to determining the effect of the shutter on darks by keeping the shutter closed and open respectively. Figure 13 shows example images of the focal plane from both of these runs. Figure 14 shows the difference between the shutter open and closed and how this condition affects the dark current. No easily visible difference can




Figure 12: (left) The original impact of the AC light leak on a 900 s dark difference image (AC on minus AC off), we observe in particular a “bright triangle” on the right of this focal plane image. We also note the presence of the persistence of ~ 10 ADU for a few sensors, the same e2v sensors than the one visible in the left figure 48. We did not collect the exactly same condition image, a dark without the AC leak mitigation with the persistence mitigation voltages (see Sect. 5.1.1). (right) The result after masking the LEDs of the motor encoders in the AC. No light associated with the FES is present in 900 s dark difference image.

be seen between the two images. This is most likely due to the shrouding of the camera (see Section 3.3.1) as well as the pinhole filter being in place for these runs.

3.5 Impact of filters on dark images

To investigate how the filter affects the dark images, we took 900 s darks with the available filters in the filter wheel: E1114 (empty filter), E1115 (*g*), E1116 (*y*), and E1117 (*r*). The heat maps of the dark currents from eo_pipe can be found in Figure 15. The primary effect of including the filters was reducing the glow due to the AC (see Fig. 12). The global average of the median amplifier dark currents decreased from 26×10^{-3} e-/s with the empty filter to 3.5×10^{-3} e-/s for *r*, 1.1×10^{-3} e-/s for *y*, and 0.63×10^{-3} e-/s for *g*. The discrepancy between the filters could arise if the AC light shines more brightly in the redder wavelengths and even the IR. Unfortunately, we were not able to obtain data with the other three filters to confirm this.

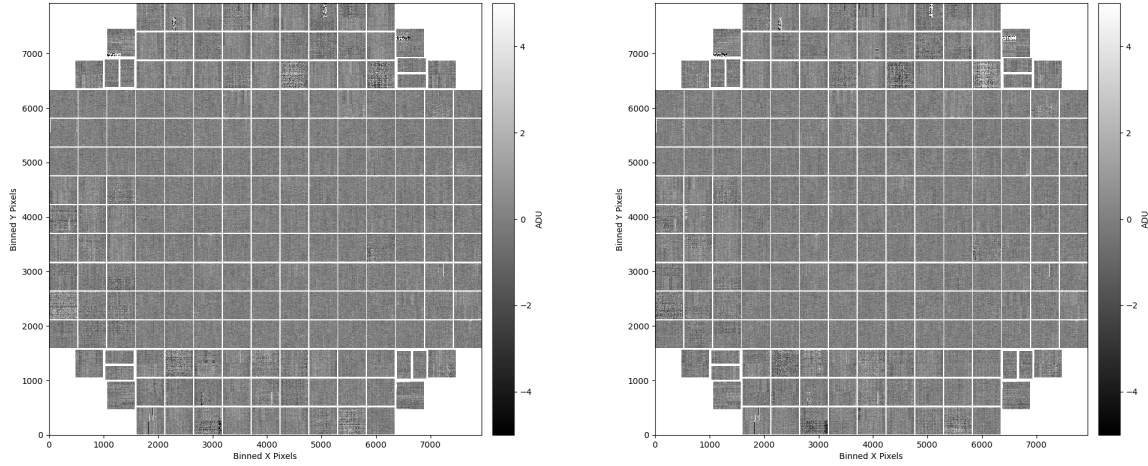


Figure 13: (left) An example image of run E1075 (shutter closed). (right) An example image of run E1076 (shutter open). If there was a light leak, we would expect to see the 21 spots of the pinhole filter. However, there is no noticeable difference between the images, confirming that our shrouding described in Section 3.3.1 is good and that the shutter being open and closed does not affect the dark current.

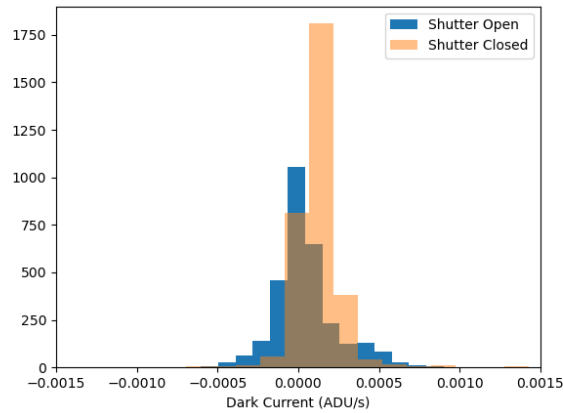


Figure 14: Histogram comparison between the dark current between the shutter being open versus closed.




Figure 15: The heat map of the dark current with the empty filter installed (E1114; top left), the g filter installed (E1115; top right), the y filter installed (E1116; bottom left), and the r filter installed (E1117; bottom right)

4 Reverification

All EO camera test data is processed through the calibration products and electro-optical pipelines to extract key metrics from the data run. The key LSST Camera metrics from Run 7, and their comparison to previous runs are discussed below.

Among the motivations for these measurements, the primary concern is whether LSST Camera has maintained its performance characteristics between Run 6 and Run 7, since LSST Camera was transported from SLAC to Cerro Pachon. The testing condition is supposed to be identical; however as described in Section 4.4.5, two rafts have slightly different voltages between two runs.

4.1 Background

Initial characterization studies performed on LSST Camera during Run 7 primarily used two image acquisition sequences.

- B protocols: this acquisition sequence consists of the minimal set of camera acquisitions for EO testing, including
 - Bias images
 - Dark images
 - Flat pairs - flat illumination images (flats) taken at varying flux levels
 - Stability flats - flats taken at constant flux levels
 - Wavelength flats - flats taken with different LEDs
 - A persistence dataset - a saturated flat, followed by several darks
- PTCs (photon transfer curves): this acquisition sequence consists of a sequence of flat pairs taken at different flux levels. The flat acquisition sequence samples different flux levels at a higher density than the B protocol flat sequence, enabling more precise estimates of flat pair metrics including pixel covariances (see Fig. 16).

For comparisons between Cerro Pachon EO runs and the final SLAC IR2 equivalents, the following runs are used (see Table 2).

The naming of the EO runs was established during initial LSST Camera integration and testing.




Figure 16: Flat-pair comparison between PTC and B protocol

Table 2: Reference runs for Run 6 and Run 7 comparisons

Run Type	Run 6	Run 7
B Protocol	13550	E1071
PTC	13591	E749

The final SLAC IR2 run from November 2023 was named “Run 6”, while the data acquisitions from Cerro Pachon from September through December 2024 are considered “Run 7”. Additionally, individual EO acquisitions are tagged with a run identifier. This is commonly referred to as a Run ID. For all SLAC runs, the run identifier was a five digit numeric code, while the Cerro Pachon runs were “E-numbers” that started with a capital E followed by a numeric code.

4.2 Stability flat metrics

4.2.1 Charge transfer inefficiency

CTI, or charge transfer inefficiency, measures the fraction of charge that fails to transfer from row to row during readout, and appears as trailing charge in the image area. Consequences of high CTI include loss of charge, distorted signals in the direction of parallel transfer, and reduced sensitivity in low light imaging. CTI measurements are made using the EPER method (Snyder et al., 2021), for which the ratio of the residual charge in the overscan pixels to the total signal charge in the imaging region is evaluated. In the context of LSST Camera, we measure

CTI along both the serial and parallel directions.

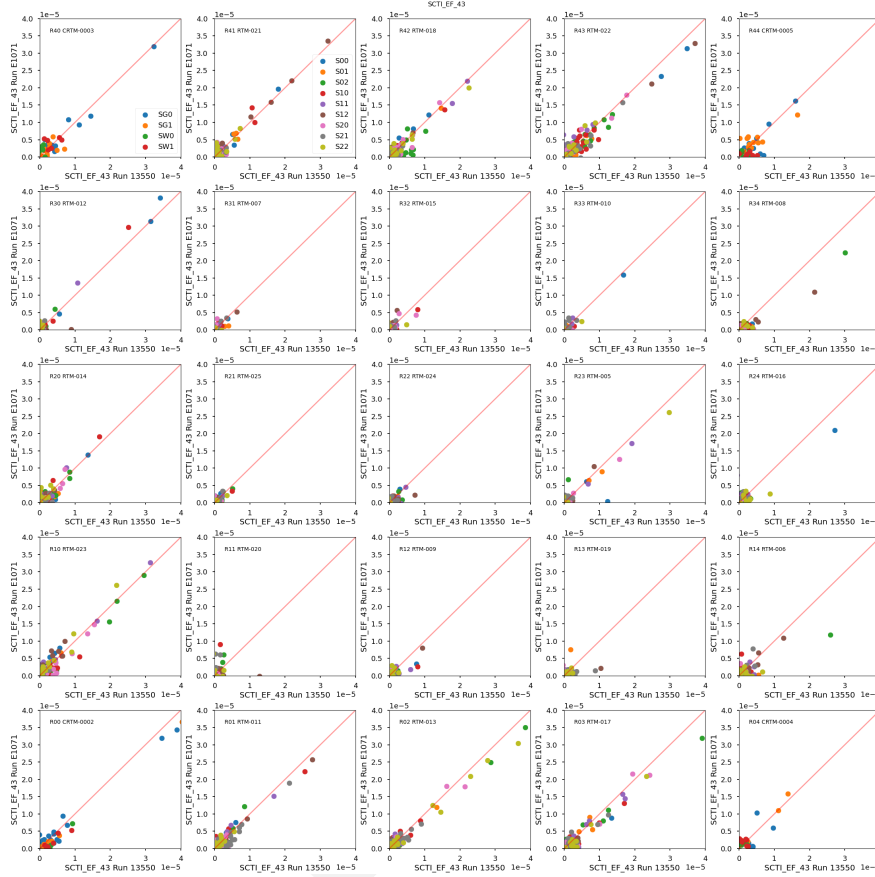


Figure 17: Serial CTI amplifier measurements separated by raft for Run 7 (E1071) and Run 6 (13550)

4.2.1.1 Serial CTI The CTI along the serial registers of the amplifier segments of the LSST Camera CCDs is consistent between Run 6 and Run 7 (Fig. 17). Both sensor types show low CTI, span a range of $\sim 2 \times 10^{-5}$ % for e2v sensors, and by $\sim 4 \times 10^{-6}$ % for ITL sensors (Fig. 18).




Figure 18: Distributions of differences in serial charge transfer inefficiencies between Run 7 (E1071) and Run 6 (13550), grouped by CCD type.

4.2.1.2 Parallel CTI The CTI along the parallel direction is consistent between Run 6 and Run 7 (Fig. 19). Both sensor types are found to have extremely low CTI on the order of 10^{-5} %, and span a range of $\sim 1 \times 10^{-5}$ % for e2v sensors, and by $\sim 7 \times 10^{-4}$ % for ITL sensors (Fig. 20). Both of these measurements pass the CTI requirements (see table 3).




Figure 19: Parallel CTI comparison by raft for Run 7 (E1017) and Run 6 (13550).




Figure 20: Distributions of differences in parallel charge transfer inefficiencies between Run 7 (E1071) and Run 6 (13550), grouped by CCD type.

4.3 Dark metrics

4.3.1 Dark current

Dark current is the small amount of electrical charge generated in the absence of light due to thermal activity within the semiconductor material of a CCD. This effect occurs when electron/hole pairs are thermally released into the conduction band in the CCD, mimicking the signal that light would produce. Dark current increases with temperature, so cooling the CCD is a common method to reduce it in sensitive imaging applications. Dark current introduces noise into an image, particularly in low-sky background conditions in long exposures. The measurement of dark includes the dark current and stray light, making them impossible to distinguish each other since they both linearly evolve with time. In the context of LSST Camera, we measure dark current from the combined dark images across all amplifiers as the upper limit.

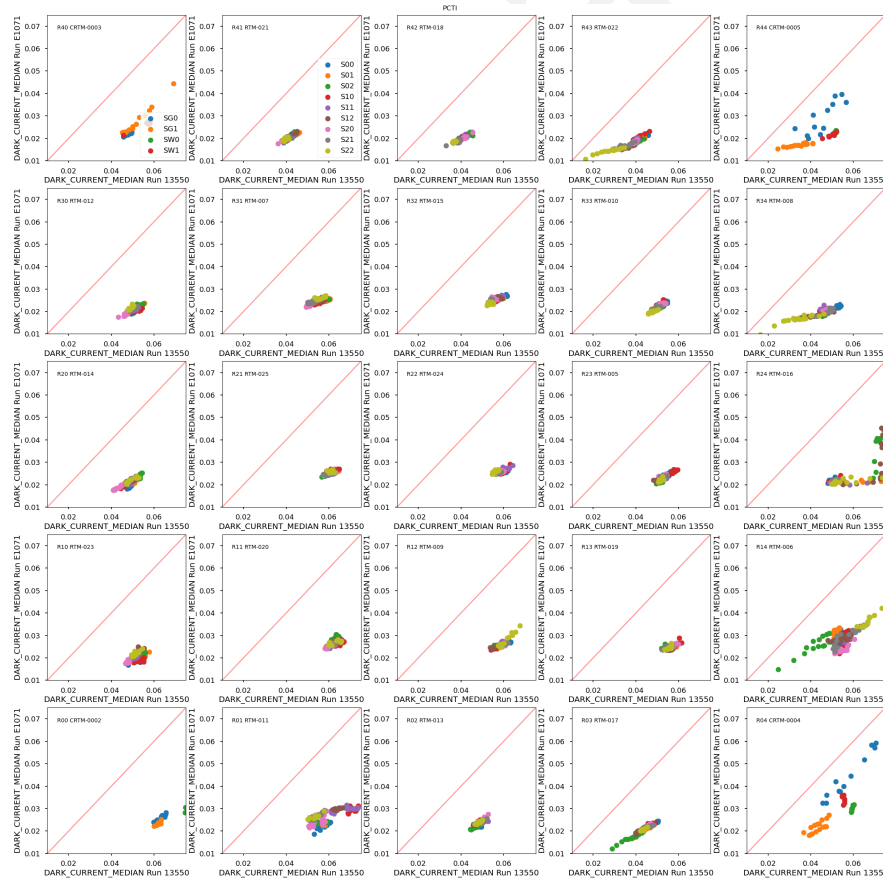


Figure 21: Dark current comparison by raft for Run 7 (E1071) and Run 6 (13550).

Unexpectedly, the dark current was significantly less in Run 7 than Run 6 (Fig. 21). We do not attach particular significance to the finding because this could be the result of improved shrouding on the camera in the Level 3 white room relative to the IR2 clean room SLAC.

Draft

4.3.2 Bright defects

Bright defects are localized regions or individual pixels that produce abnormally high signal levels, even in the absence of light. These defects are typically caused by imperfections in the semiconductor material or manufacturing process of the CCD. Bright defects can manifest as “hot pixels” with consistently high dark current, small clusters of pixels with elevated dark current, or as “hot columns” (pixels along the same column that have high dark current).

In the context of LSST Camera, we identify and exclude bright pixels from the dark current measurement, with the threshold for a bright defect set at $5 \text{ e}^-/\text{pix/s}$, above which the pixel-cluster/column is registered as a bright defect. In addition to the bright pixel metric, eo-pipe also computes a bright column metric, which is any region of bright pixels that is contiguous over 50 pixels or more.

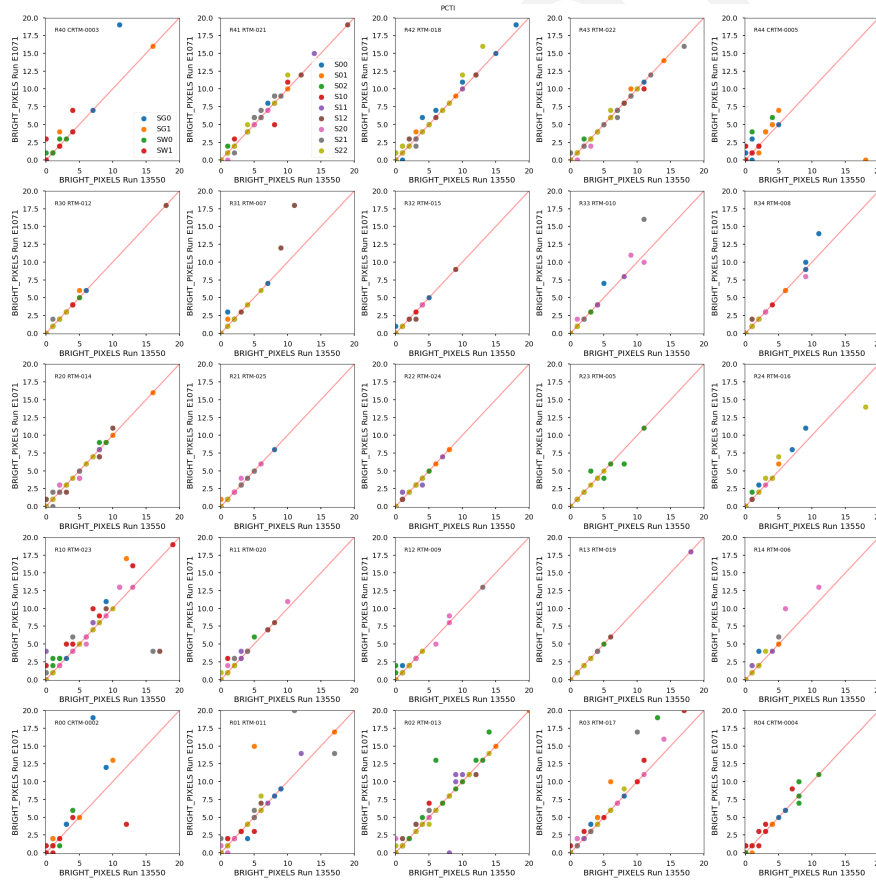


Figure 22: Bright pixel comparison by raft for Run 7 (E1071) and Run 6 (13550)

Evaluating the change in defect counts on each amplifier segment between Run 6 and Run 7,

and aggregating the amplifiers by the detector manufacturer shows a small increase of bright defects in Run 7 (Fig. 22). Figure 23) displays differences of the measurements. The median values agree well, while there are signs of the positive tail. For ITL sensors, we find that 12% of the amplifiers have more bright pixels than in Run 6. For e2v sensors, we find 4% of the amplifiers that have more bright pixels. Despite this, the number of bright defects between runs does not increase for most sensors.

The reason is not totally clear, but the difference in the illumination pattern as described in Section 3.1 might play a role, which implies that a small number of defects could be created in the CCOB optical path.




Figure 23: Distributions of differences in bright pixel count per amplifier between Run 7 (E1071) and Run 6 (13550), grouped by CCD type.

4.4 Flat pair metrics




Figure 24: A comparison of Run 6 and Run 7 PTCs for a central amplifier.

4.4.1 Linearity and PTC turnoff

Linearity turnoff and PTC turnoff are two closely related metrics used to characterize the upper limit of the usable signal range for accurate shape measurements and photometry. Linearity turnoff is the signal level above which the PTC curve (Figure 24) deviates from linearity and is measured for each amplifier segment of each CCD. We have defined the deviation threshold as 2%. PTC turnoff refers to the high-signal region of the PTC above which the PTC variance decreases with increasing signal. This is due to saturation within the pixel wells of the CCDs. While slightly different, both metrics provide important information about the upper limits of the dynamic range in our sensors. Linearity turnoff is measured in units of e^- , while PTC turnoff is measured in ADU.

In our linearity turnoff measurements, we find close agreement between our Run 7 and Run 6 measurements for both ITL and e2v sensors.

Run 7 PTC turnoff measurements agree closely between Run 6 and Run 7, differing by $\leq 200 e^-$ for both ITL and e2v sensors. Notably, they are lower on average for both detector types.

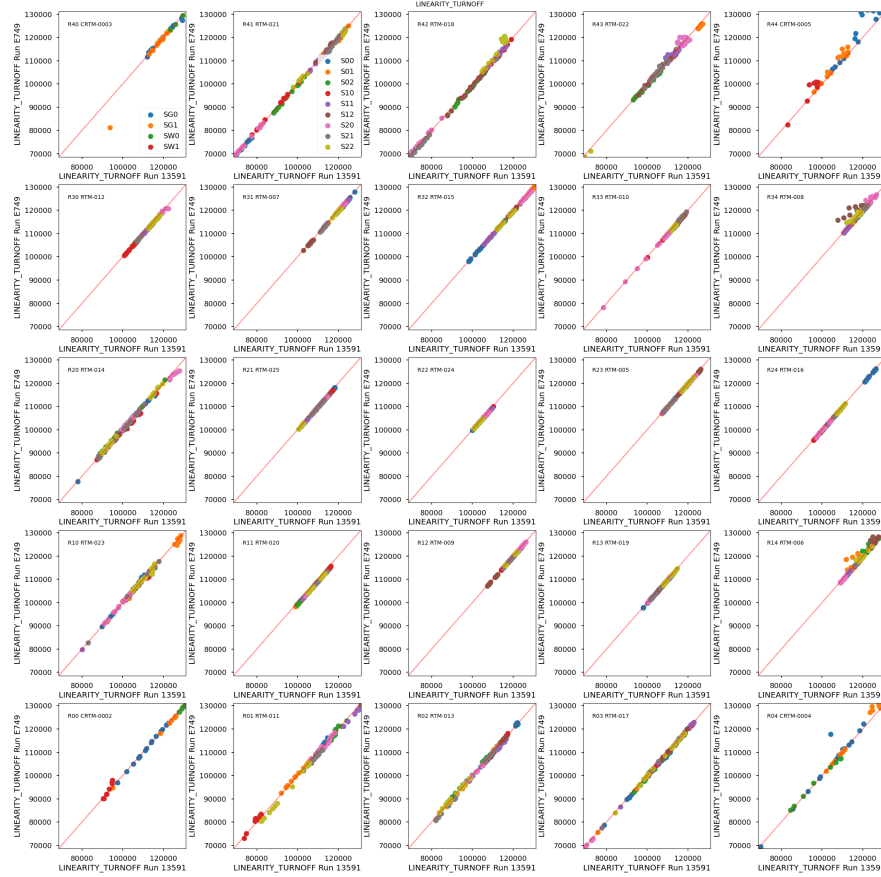


Figure 25: A comparison of Run 7 amplifier measurements of linearity turnoff, separated by sensor type. For both sensor types, measurements agree across both runs.




Figure 26: A comparison of Run 7 amplifier measurements of linearity turnoff, separated by sensor type. For both sensor types, linearity turnoff is above the 90k e- specification for a majority of amplifiers. A subset of ITL amplifiers are below the 90k e- threshold, while two e2v amplifiers are below that specification.




Figure 27: A comparison of Run 6 and Run 7 amplifier differences in PTC turnoff, separated by sensor type. For both sensor types, PTC turnoff is very consistent.

4.4.2 PTC gain

PTC gain is the conversion factor between digital output signal and the the number of electrons generated in the pixels of the CCD. It is one of the key parameters derived from the Photon Transfer Curve, as it is the slope above the flux range at which the variance is dominated by shot noise, and below the PTC turnoff. Gain is expressed in e^-/ADU , and scales the digitized analog signals from the ASPICs (Application Specific Photonic Integrated Circuits) to units of e^- .

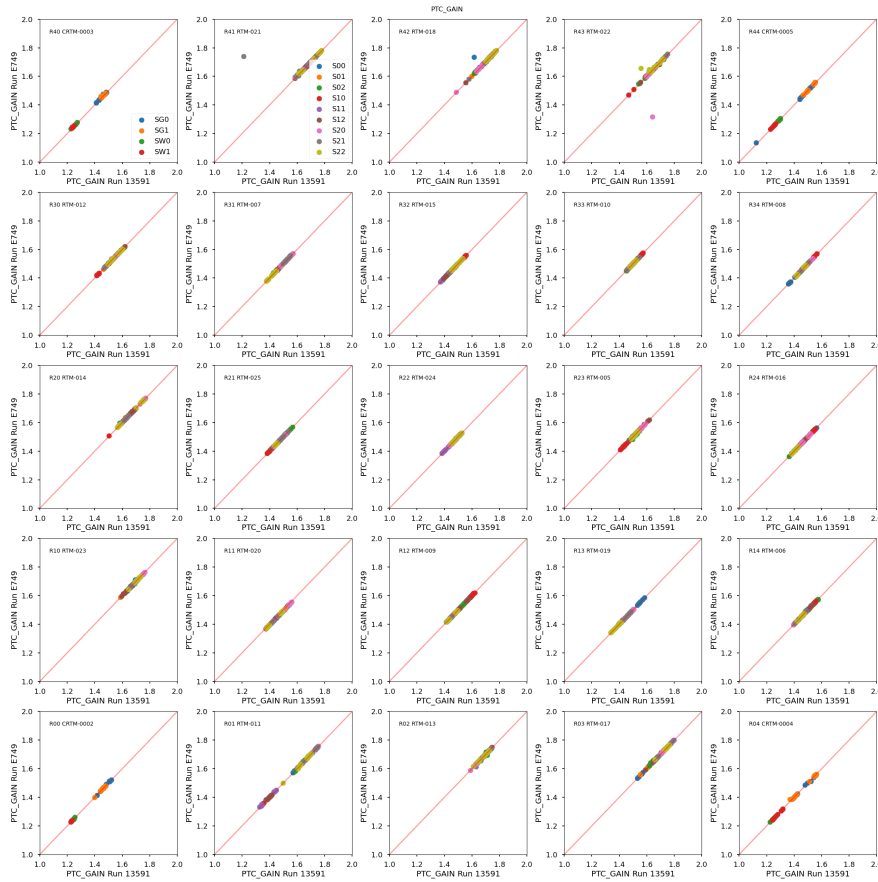


Figure 28: A comparison of Run 6 and Run 7 amplifier measurements in gain, separated by sensor type. For both sensor types, gain is very consistent.

PTC gain measurements agree extremely closely across all sensors in the focal plane.

4.4.3 Read Noise

Read noise is induced charge during the readout process of the LSST Camera sensors. Common causes of read noise include thermal noise in the electronics or imperfect charge transfer. Read noise is measured in electrons, and is calculated by taking the standard deviation of the overscan region.

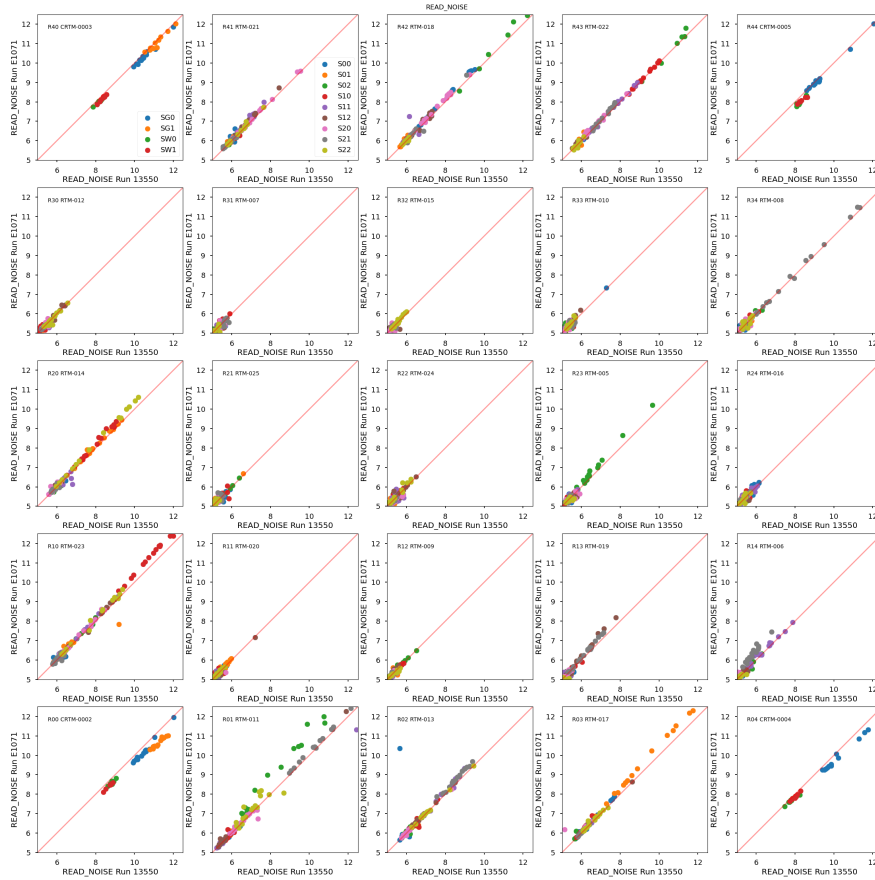


Figure 29: A comparison of Run 6 and Run 7 amplifier measurements for read noise, separated by sensor type. For both sensor types, read noise is higher in Run 7.

The read noise is higher in Run 7 than Run 6 for both sensor types. The difference is on the order of 0.05 e⁻.

4.4.4 PTC Noise

PTC noise is a fitted parameter in the PTC model, and is the foundation of the PTC model. In the shot noise regime, where the PTC slope is 1/2 in log-log space (from 500 - 10k ADU in




Figure 30: A comparison of Run 6 and Run 7 amplifier measurements for read noise, separated by sensor type. For both sensor types, read noise is higher in Run 7.

figure 24), the data can be fitted to a line. The intercept of that line is the PTC noise, measured in ADU.

PTC noise differs by $\lesssim 0.1$ e- on average for science sensors. Notably, e2v sensors measure a PTC noise 0.02 e- lower in Run 7, while ITL sensors measure a PTC noise 0.08 e- higher in Run 7 (see figure 32).

4.4.5 Brighter-fatter coefficients

The brighter-fatter effect in CCDs refers to the phenomenon where brighter sources appear larger (or “fatter” than dimmer ones). This occurs due to electrostatic interactions within the pixel wells of the CCDs, when a pixel accumulates a high charge from incoming photons and creates an electric field that slightly repels incoming charge carriers into neighboring pixels.

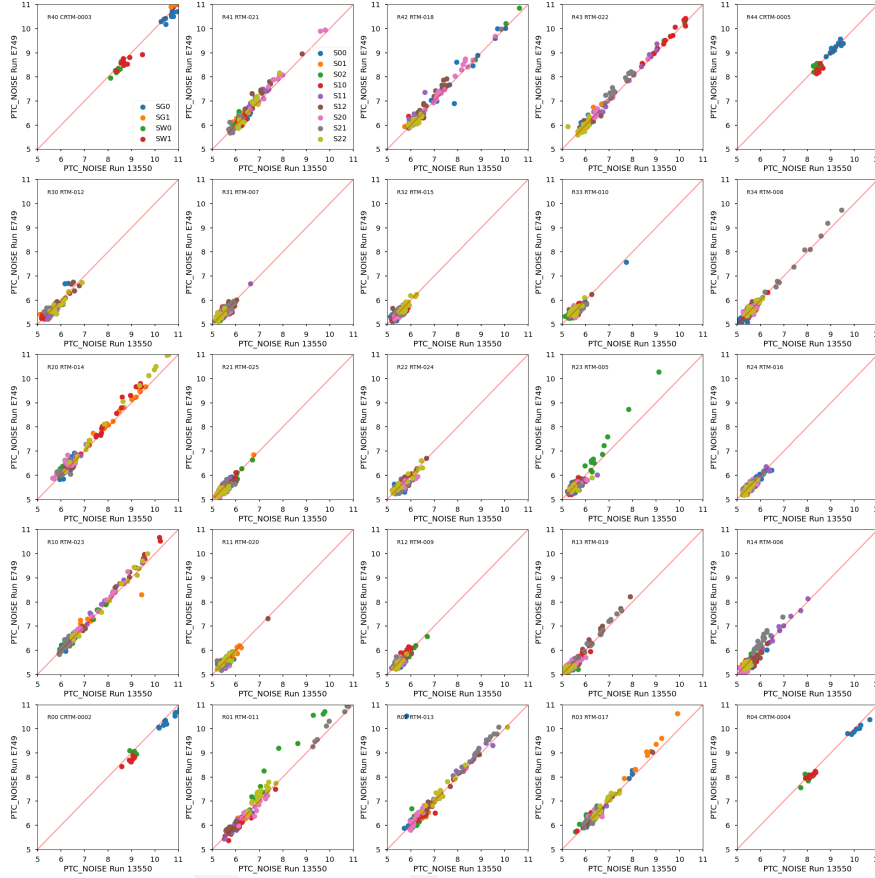


Figure 31: Comparison of amplifier measurements of PTC noise for initial and final Run 7 conditions.

The brighter-fatter effect can be modeled as the most dominant source of pixel-pixel correlations. Following the PTC model from Astier et al. (2019), a_{00} describes the change of a pixel area due to its own charge content, or the relative strength of the brighter-fatter effect. Since same-charge carriers repel each other, the pixel area decreases as charge accumulates inside the pixel well, which implies $a_{00} < 0$. Similarly a_{10} describes the area change caused by a pixel to its nearest serial neighbor, and a_{01} to the parallel nearest neighbor. Figs. 33 and 34 compare the measurement of these coefficients carried out at SLAC and at the summit. We see that the variations are modest (and could be explained by noise) except for two rafts: R10 and R11. The Run 6 data used for this comparison was acquired with a high voltage of 45V applied to these two rafts, rather than the usual 50V. The sensitivity of our measurements of the brighter-fatter coefficients is sufficient to detect the change of electrostatic conditions due to this change of drift field in the sensors. In `eo_pipe`, an absolute value is taken of the a_{00} parameter, so the tabulated quantities are positive.

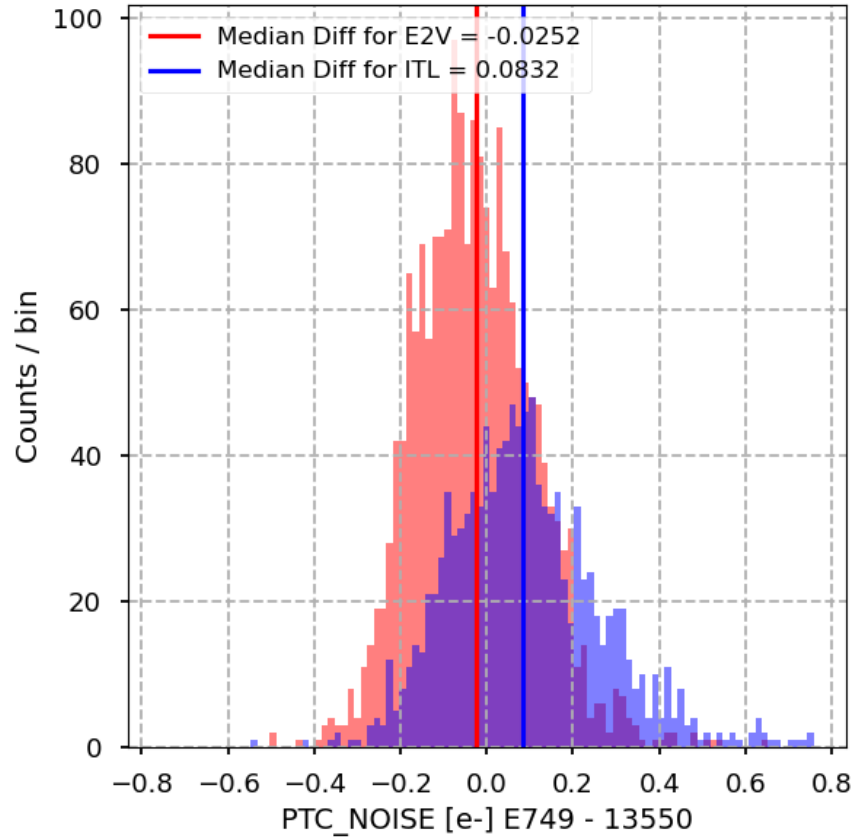


Figure 32: Comparison of the PTC noise measurements in different sensors and run types.

The distribution of the difference of a_{00} measurements between the runs is displayed in Figures 36 and 35.

However, the differences in the brighter-fatter a_{00} coefficient between Run 6 and Run 7 show that the magnitude of a_{00} decreased for most of the outliers, which implies an improvement in imaging for those pixels.

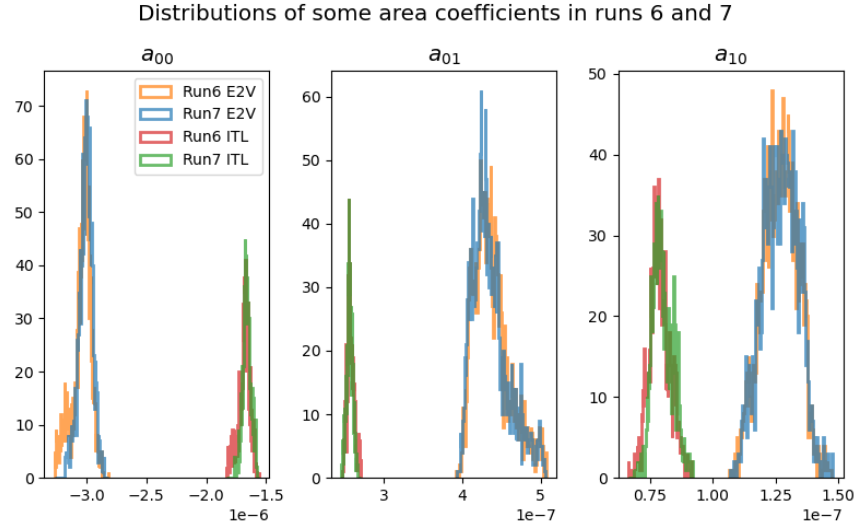


Figure 33: Distributions of a_{00} , a_{01} and a_{10} in Run 6 and Run 7. Those are very similar, except for two rafts, which had a lower drift electric field (45 V vs. 50 V) in Run 6. The ratios are displayed in Figure 34.




Figure 34: Ratio of measurements of a_{00} , a_{01} and a_{10} coefficients (one per amplifier) for Run 6 and Run 7. They are very consistent, except for two rafts (R10 and R11) where the high voltage was changed between the two runs.




Figure 35: Comparison of Run 6 and Run 7 amplifier differences in the a_{00} coefficient, separated by amplifier across the focal plane. The amplifiers associated with R10 and R11 are noted outliers, as seen in figure 34.




Figure 36: A comparison of Run 6 and Run 7 amplifier differences in the a_{00} coefficient, separated by sensor type. For both sensor types, the a_{00} coefficient is very consistent. The two peaks on the left represent the two outlier rafts visible on Figure 34. The a_{00} values are of the order of 2 to 3 $10^{-6} e^{-1}$.

4.4.6 Row-means variance

Row-means variance is a metric that measures the mean row-to-row variance of differences between a pair of flats. By computing variance of means of differenced rows at each flux level, we can measure any changes in gain row-by-row and also changes in correlated noise along with row.




Figure 37: A comparison of Run 6 and Run 7 amplifier differences in row-mean-variance slope. For both sensor types, row-means-variance slope is weaker in Run 7. This is more pronounced for e2v sensors.

Differences in row-means variance between runs are evident, and are distinctly different for different detector types. The difference between runs is more significant for ITL sensors, ~9% smaller on average in Run 7. For e2v sensors, the effect is ~3% smaller in Run 7. This indicates that either row-by-row correlated noise or row-by-row gain change is less in Run 6. Since we did not change the sequencer file, the most natural explanation is the row-by-row correlated noise. But further investigation is needed.




Figure 38: A comparison of Run 6 and Run 7 amplifier differences in row-mean-variance slope, separated by sensor type. For both sensor types, row-means-variance slope is weaker in Run 7. This is more pronounced for e2v sensors.

Divisadero tearing (or Rabbit ears) manifests as signal variations near amplifier boundaries, connected features that are often jagged Juramy et al. (2020); Utsumi et al. (2024). These variations are on the order of $\sim 1\%$ relative to the flat field signal. To quantify divisadero tearing in a given column, we measure the column signal, and compare it to the mean column signal from flat fields.




Figure 39: A comparison of Run 6 and Run 7 amplifier differences in Divisadero tearing, separated by sensor type. For both sensor types, divisadero tearing is weaker in Run 7. The difference is more pronounced for e2v sensors, which have larger divisadero tearing in general.

Divisadero tearing is broadly consistent between Run 6 and Run 7, with both sensor types demonstrating lower divisadero tearing in Run 7. Taking amplifier differences, e2v sensors show a weaker divisadero signal in Run 7 by 0.1%, while ITL sensors demonstrate a weaker Divisadero signal in Run 7 by 0.05% (see Fig. 40).




Figure 40: Comparison of Run 6 and Run 7 amplifier differences in divisadero tearing, separated by sensor type. For both sensor types, divisadero tearing is weaker in Run 7.

4.4.8 Dark defects

Dark defects are localized regions or individual pixels that produce abnormally low signal levels, even in the presence of light. Similarly to bright pixels, dark pixels are also quantified in dark columns over 50 pixel contiguous regions. These defects are caused by imperfections in the semiconductor material, imperfections during the manufacturing process of a CCD. For our evaluation, we extract dark pixels from combined flats, with the threshold for a dark defect defined as a -20% deficit from the average flat field flux measured in the image segment.




Figure 41: Illustration of masked border pixels (yellow) for detector 85 (R21_S11). The average defect mask size is 4 pixels along the serial (x-pixel) direction, and 5 pixels along the parallel direction. Additional dark defects exist in the sensor, but are difficult to quantify due to the overwhelming contribution from the picture frame response.

The eo-pipe configuration for evaluating dark defects considers a border pixel region that is masked differently from the dark pixels. The default size for this edge is zero pixels. With a zero pixel border mask, the average dark defect count is 1800 per amplifier, with $\geq 95\%$ of the contribution coming from the picture frame. The ‘picture-frame response’ (also called ‘edge roll-off’) near the edges of the sensors is due to a decrease in the pixel active area. It is difficult to extract useful information about the dark defects in the focal plane without excluding the picture frame. The effects of the picture frame signal on dark defect masking is shown in figure 41.

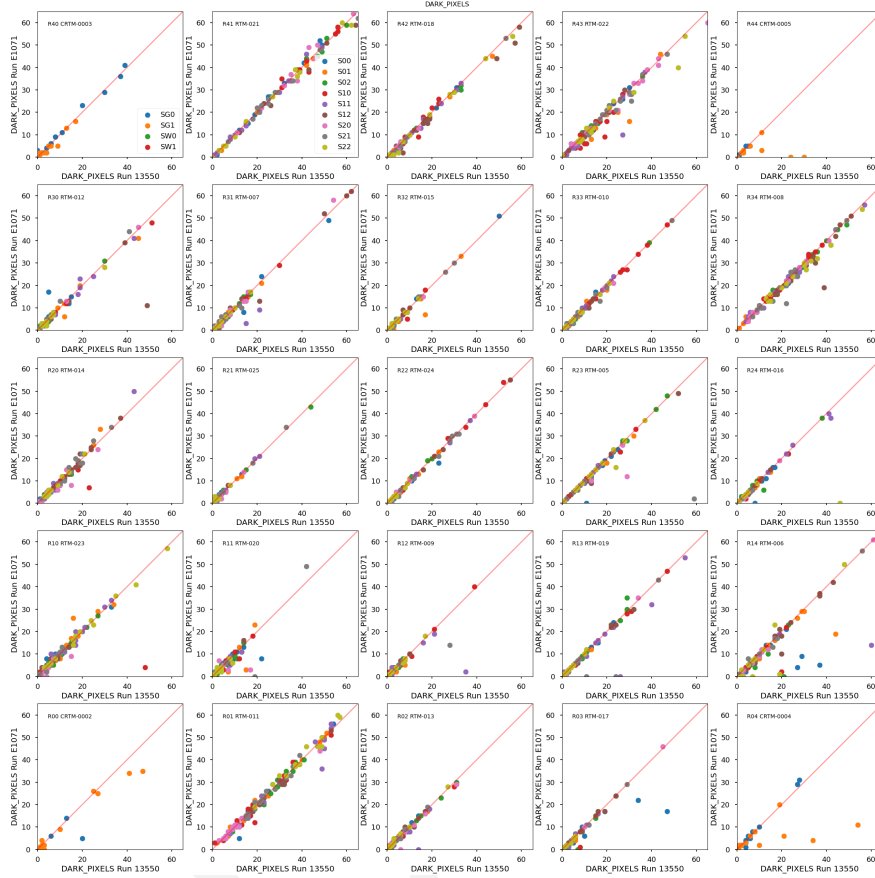


Figure 42: Comparison of dark pixel counts in Run 7 (E1071) and Run 6 (13550), with separate plots for each raft. Within each plot the color coding for all amplifier segments in a given CCD is the same.

The default eo-pipe configuration has no border masking. The largest region permitted for the picture frame region is 9 pixels, determined by LCA-19636. Using a 9 pixel mask, the picture frame signal is removed, leaving true dark defects to be measured without contamination.

In both instances, the contamination of dark pixels across the focal plane is ≤ 10 pixels per amplifier on average. There is a measurable improvement in the dark pixel counts, decreasing by one pixel per amplifier between Run 6 and Run 7.




Figure 43: Comparison of dark pixel counts in Run 7 (E1071) and Run 6 (13550). Top: A histogram of amplifier measurements, separated by run number and sensor type. Bottom: A histogram of the amplifier dark pixel count differences, the difference is taken as the measurement from Run 6 and the measurement from Run 7.

4.5 Persistence

Persistence is a feature of CCDs and how they are operated involving charge trapped in the surface layer after high-flux exposures (Banovetz et al., DMTN-276; Utsumi et al., 2024). Persistence is described in detail in Section 5.1. Here we consider the measurements taken as part of a persistence measurement task in the typical B protocol. For measuring persistence, a high-flux acquisition is taken, followed by a sequence of dark images. The persistence signal has been observed to decrease in subsequent dark images as the trapped charge is released (see Figure 44 for an example). As a metric for persistence, we evaluate the difference between the residual ADU in the first dark image and the average of the residual ADU in the final dark images. This residual signal is found to be ~10 ADU.




Figure 44: Persistence signal observed in R22_S11 in Run 7 (E1110) as a function of time after the high-flux flat image. The color coding indicates the individual amplifier segments. The persistence metric is defined as the residual signal in the first dark image after the flat acquisition (red box). Note that over time the signal does not decay entirely to zero. This may be more due to bias fluctuation or incomplete image reduction. It definitely should return to 0 on some timescale.

In the initial Run 7 measurements, we had not changed any operating parameters of LSST Camera, so we would expect persistence to still be present images at the same level as in Run 6.

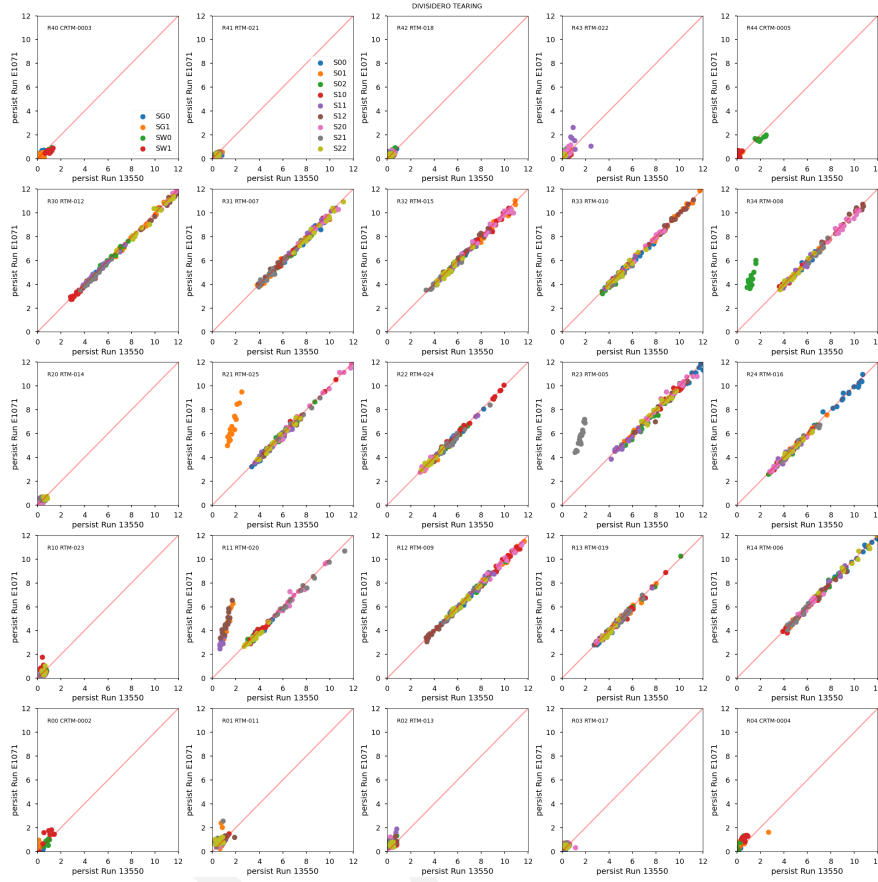


Figure 45: Comparison of persistence metric between Run 7 (E1071) and Run 6 (13350), organized by raft. The color coding indicates individual CCDs. Several e2v CCDs have markedly greater persistence in Run 7.

The persistence signal is generally consistent in e2v sensors between Run 6 and Run 7. Several e2v CCDs have greater persistence metric value in Run 7 (Fig. 45). The outliers in persistence measurements are due to higher initial residual ADU in a subset of rafts, resulting in an excess of ~5 ADU when comparing Run 6 with Run 7 (see Fig. 46).




Figure 46: Comparison of persistence profiles for R12_S21 between (left) Run 6 (13550) and (right) Run 7 (E1071). The decay time constants are similar but the initial persistence level is greater in Run 7. The asymptotic levels are also slightly different.

4.6 Differences between Run 6 and Run 7

All camera performance metrics from the summit show close agreement with SLAC IR2 tests. PTC/full-well metrics were consistent, and no significant bright cosmetic defects developed. Dark cosmetic defects are difficult to quantify due to the edge sensor effects, though the consistency in CTI measurements would indicate that dark defects did not change from previous runs. Dark current and divisadero tearing show improved performance compared to Run 6, while the Persistence feature is still prominent in e2v sensors.

Parameter [unit]	Specification	e2v		ITL	
		Run 6	Run 7	Run 6	Run 7
Serial CTI [%]	5×10^{-4}	3.7×10^{-5}	1.1×10^{-5}	1.2×10^{-4}	1.7×10^{-4}
Parallel CTI [%]	3×10^{-4}	1.2×10^{-5}	1.2×10^{-5}	3.4×10^{-7}	-4.8×10^{-6}
Dark current [e-/pix/s]	None	0.055	0.025	0.046	0.021
Bright defects [count]	None	0	0	0	0
Linearity turnoff [e-]	>90,000 e-	156,000	168,000	173,000	178,000
PTC turnoff [e-]	>90,000 e-	126,000	133,000	117,000	129,000
PTC Gain [e- / ADU]	None	1.48	1.48	1.67	1.68
Read noise [e-]	<9 e-	5.30	5.32	6.20	6.26
PTC $a_{00} [\frac{1}{\text{pix}^2}]$	None	3.09×10^{-6}	3.09×10^{-6}	1.71×10^{-6}	1.70×10^{-6}
BF x-correlation	None	0.524	0.517	0.716	0.752
BF y-correlation	None	0.179	0.170	0.286	0.287
Row-means variance	None	0.993	0.884	0.992	0.947
Dark defects [count]	<2%	4	3	9	8
Divisadero tearing maximum [%]	None	0.327	0.274	0.752	0.626
Persistence [ADU]	None	5.67	5.64	0.48	0.42

Table 3: Comparison of the median values of different parameters between Run 6 and Run 7, separated by detector type. For this comparison, only science detectors are considered.

5 Camera Optimization

5.1 Persistence optimization

Leftover signal (“persistence”) in the first dark image acquired after intense illumination has been observed (Sec. 4.5). Persistence has been observed in an early prototype e2v sensor as early as 2014 (Doherty et al., 2014). It was confirmed that the amplitude of the persistence decreased as the parallel swing voltage was decreased. This is consistent with the persistence being a Residual Surface Image (RSI) effect as described in Janesick (2001), i.e., the excess charges are being trapped at the surface layer. The level of persistence is about 10–20 ADU, and the decay time constant is about 30 s (Banovetz et al., DMTN-276).

During the EO testing in 2021 (e.g., Run 13177 in Run 5), we also found the persistence made a streak toward the readout direction from the place where bright spot illumination occurred in a previous image. We call this “trailing persistence”.

As noted in Section 4.4.7, depending on operating conditions e2v sensors have another major non-ideality, so-called “tearing”, which is considered a consequence of the nonuniform distribution of holes. Over the past few years, our primary focus in the optimization of the operating parameters was mitigation of the tearing, and we successfully eliminated the tearing by changing the e2v voltages from unipolar (both parallel rails high and low are kept at positive voltages) to bipolar (the parallel high is positive, and the low is negative) following the Bipolar voltage formula. However, the persistence issue remained unchanged.

If the persistence is an RSI, two approaches could be taken as discussed in Utsumi et al. (2024): either 1) establishing the pinning condition where the holes make a thin layer at the front surface so that the excess charges recombine with the holes, or 2) narrowing the parallel swing so that the accumulated charges in the silicon do not get close to the surface state.

The pinning condition could be established by decreasing the parallel low voltage to as low as -7.0 V or lower. The transition voltage needs to be empirically determined. However, Teledyne e2v advised that the measured current flow increases as the parallel low voltage is decreased, which increases the risk of damaging the sensor by inducing a breakdown¹. Also, the excess charges could be recombined by the thin layer of holes, which could affect linearity at high

¹We note that ITL operates at a parallel low voltage of -8.0 V. We have observed the increased current flow. But we have software protection so that the current does not increase too much.

flux levels when charges start to interact with the holes.

The parallel swing determines the full-well capacity. Depending on whether the accumulated charges spread over the columns or interact with the surface layer, blooming full-well regime and surface full-well regimes can become relevant. A full-well level between these two regimes is considered to be optimal (Janesick, 2001), with no persistence and dynamic range as great as possible. Because we observe the persistence effect, we likely operate the sensor in the surface full-well condition and we need to decrease the parallel swing to get the blooming full-well or the optimal full-well level. The obvious downside is decreasing the full-well capacity.

The sensor control voltages are defined relative to each other. Changing, e.g., the parallel swing also requires changes to all other voltages to operate the sensor properly, e.g., to properly reset the amplifier. The initial voltages were given in the original Bipolar formula but to decrease the parallel swing we had to switch to the new persistence mitigation formula in order to satisfy the constraints (Persistence mitigation voltage).

Snyder et al. (2024), set up a single sensor test-stand at UC Davis. They attempted multiple different approaches mentioned above and reported the results in Polin et al. (2025). The summary is as follows:

- The new voltages following the persistence mitigation voltage rule produce visually reasonable bias, dark, flat images.
- Narrowing the parallel swing eliminates the persistence.
- Just decreasing the collecting voltage (P_{clkH}) even without changing the swing will decrease persistence as the charge packet is not pulled as close to the surface.
- Lowering the parallel low voltage did not work as we expected; going to a more negative voltage is probably needed.

Note that the e2v sensor in the UCD setup did not exhibit the same magnitude of persistence as the e2v sensors in LSSTCam. This might be due to the characteristics of the sensor, or perhaps the differences in the electronics (e.g., the long cable between CCD and REB in the UCD setup). At UCD both the parallel high and low voltages need to be increased to reproduce persistence at a level similar to that observed in LSSTCam.

5.1.1 Persistence optimization

Based on this test result, we decided to test the new voltages with the narrower parallel swing on the LSSTCam focal plane. Keeping the parallel low voltage at -6.0V in order to operate the sensor safely (very conservative limit), we changed the parallel swing voltage from 9.3V to 8.0V as well as all the other voltages using the new formula. We overexposed the CCDs and took 20 darks afterward. Figure 47 compares the mean and median of pixel-by-pixel differences between the first and the last dark exposures, as a function of the parallel swing. (We note that this is not the persistence metric defined in Sec. 4.5, but almost identical.) As the parallel swing is decreased, the residual signal decreases, reaching roughly 10× less than the original level at 9.3V. Although we sampled at 8.0 (E1363), 8.4 (E1430), 8.65 (E1411), 8.8 (E1424) and 9.3V (E1110), 8.0V appears to work the best and could be lower with the penalty of further decreasing the full-well capacity.




Figure 47: The remaining charges measured in every e2v amplifier but aggregated by mean and median as a function of the parallel clock swing are shown.

Figure 48 displays how the persistence is reduced by the parallel swing decrease. The images were processed with the standard instrumental signature removal and assembled in the full

focal-plane view. The dark exposure was taken immediately after a 400 ke-equivalent flat exposure. The left-hand image shows the distinct pattern of elevated signal associated with the e2v sensors, which fill the inner part of the focal plane.

The right-hand image shows the same dark exposure but taken with the narrow parallel swing voltage of 8.0 V. The distinct pattern of persistence in the e2v sensors is mitigated to the (low) level of the ITL sensors.

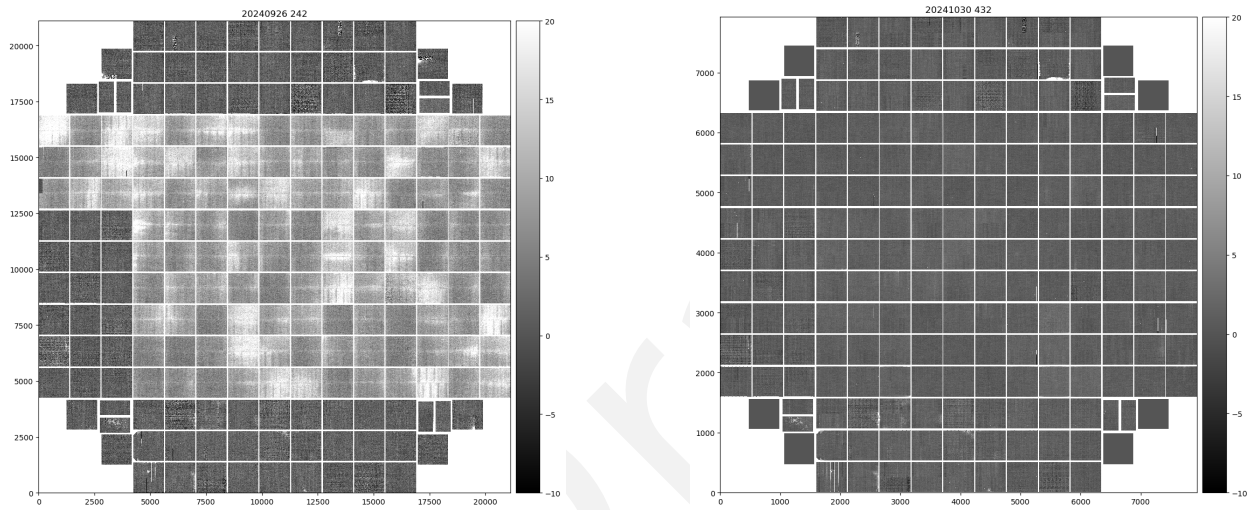


Figure 48: Comparison of dark exposures under different parallel swings. (left) The first dark exposure after a 400 ke⁻ flat image under the parallel swing of 9.3 V (Run E1110); (right) The first dark exposure after a 400 ke⁻ flat image under the parallel swing of 8.0 V (Run E1880). The figure shows no distinct patterns from persistence in e2v sensors. Note that the guide sensors were not displayed here because they were being operated in guider mode. Also some of the residuals caused by defects in ITL sensors disappeared here because of the employment of the new sequencer file (v30).

5.1.2 Impact on full-well

Reduction of the full well is expected from narrowing the parallel swing voltage. This subsection explores how much reduction in the PTC turnoff level is observed in the dense PTC runs. Two runs were acquired with identical settings except for the CCD operating voltage (E1113 for 9.3 V and E1335 for 8.0 V). As the PTC turnoff is defined in ADU, it needs to be multiplied by PTC_GAIN to compare the turnoff values in electrons. Figure 49 compares the PTC turnoffs in electrons and also shows their fractional difference. The medians of the peaks are 133,065e⁻ and 102,728e⁻, and the median reduction was 22%.




Figure 49: Histograms of the PTC turnoff values scaled to electron units (left) and the ratios of differences (right) between E1113 (9.3 V) vs. E1335 (8.0 V). The reduction of the median is 23%.

5.1.3 Impact on brighter-fatter effect

Reducing the parallel swing is expected to enhance the brighter-fatter effect (BFE), possibly in an anisotropic way. The BFE can be characterized via the evolution of the variance and covariances of flat field exposures as a function of flux, i.e., via a PTC analysis. To evaluate the impact of reducing the parallel voltage swing on e2v sensors, we acquired two series of flat field exposures with the respective voltage setups and extracted the “area” coefficients a_{ij} (Equation (1) in Astier & Regnault, 2023). The area coefficients describe how much a unit charge stored in a pixel will alter the area of some other pixel (or itself). We find that reducing the parallel swing from 9.3 V to 8.0 V typically increases the area coefficients by 10% (between 5 and 19% depending on distance indexed by i, j), and the increase is almost isotropic (i.e., very similar along serial and parallel directions; see Fig. 50). From these measurements, we anticipate that the increase of star sizes with flux in LSST data will not become more anisotropic at 8.0 V than

it was at 9.3 V, and hence this reduction of parallel swing does not risk increasing systematic uncertainty of the PSF ellipticity.

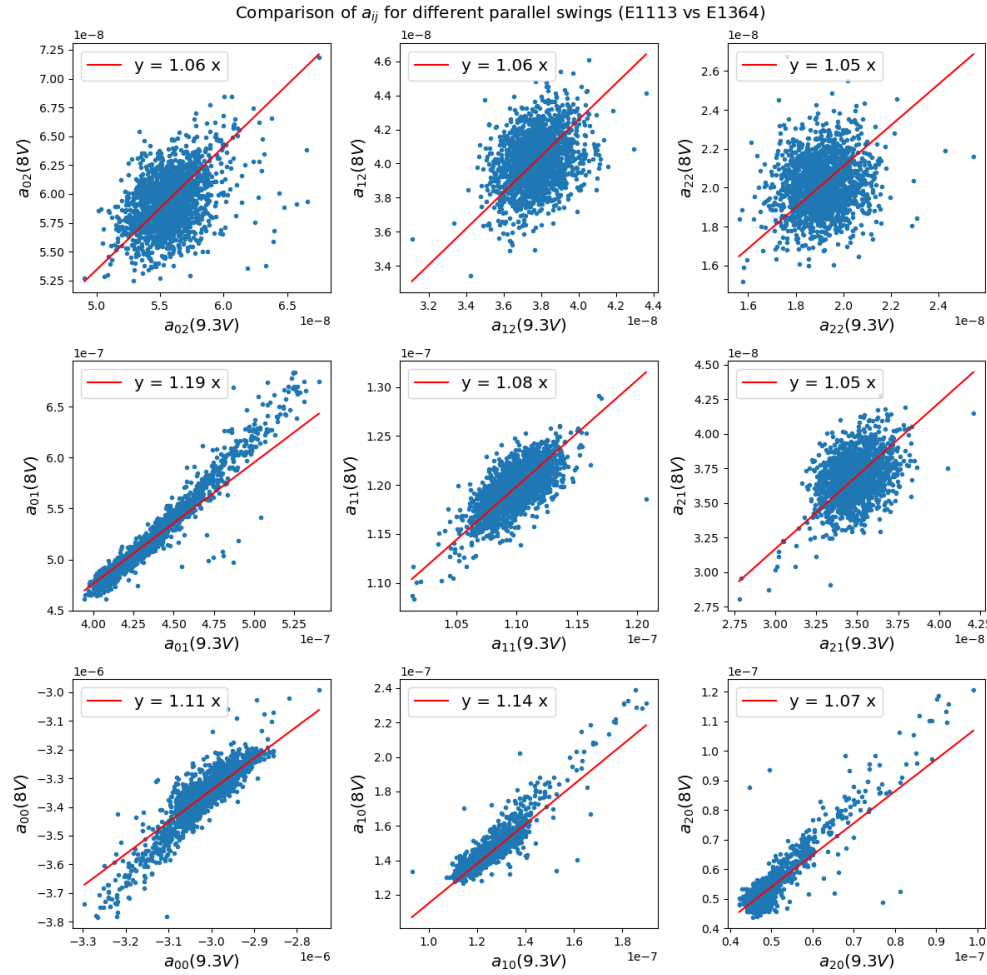


Figure 50: Scatter plots of area coefficients a_{ij} (one entry per amplifier) measured at 8.0V and 9.3V. The sub-figures correspond to separations in rows (j) and columns (i) between the source of the area distortion and its target, with the self-interaction coefficient a_{00} at the bottom left. The first neighbors increase respectively by 19% in the parallel direction and by 14% in the serial direction. So the BFE is slightly greater at 8.0V but not dramatically more anisotropic: the ratio of parallel to serial nearest-neighbor correlations increases only from 3.43 to 3.54 with the reduction of the parallel swing.

5.2 Sequencer optimization

Several efforts were undertaken to optimize the sequencer configurations during Run 7. The following points summarize the key investigations:

- **Clear:** Addressing the leftover charges at the image/serial register, discussed in Section 5.2.1.
- **Whether toggling the RG output during the parallel transfer for the e2v sensors is needed:** As such configuration flatten the bias shape for ITL device, the same type of configuration has been tested for e2v sensors during Run 7. Details are provided in Section 5.2.2
- **Whether running IDLE_FLUSH is needed:** Whether running IDLE_FLUSH had an effect on the tearing. Details are presented in Section 5.2.3
- **Phase overlap during parallel transfer for e2v:** e2v sensors feature four parallel phases. To improve the uniformity of the full well across a sensor, overlapping two phases during each time slice of the parallel transfer was introduced.
 - Sequencer files that are based on the regular v29 but have changes in the parallel transfer by having a half overlap of what it was in the original (`_halfoverlapping.seq`), a small amount of overlap compared to what it was in the original (`_overlap113.seq`), and overlapping at all (`_nonoverlapping.seq`) are created.
 - Any overlap is known to cause trailing persistence (Polin et al., 2025). We conducted several runs using both half-overlapping (E1245) and non-overlapping (E1396) sequencers but we have not studied these because the trailing persistence is no longer a concern after optimizing the operating voltages to avoid charge trapping.

5.2.1 Improved clear

5.2.1.1 Overview In this section, we describe the work done during Run 7 to improve the image clear prior to collecting a new exposure.

The problem we wanted to address is the presence of residual charges in the first rows read for an image taken just after the clear of a saturated image. These “hard to clear” charges are associated with highly saturated flats or columns (or stars as observed in AuxTel or ComCam), which leave signal in the first rows of the subsequent exposure. The effect has a sensor-specific signature:

- For all ITL CCDs (except R01_S10 for which the effect is much more significant and which will be addressed later in this section): After a very bright exposure that saturates the

overscan, the first row of the subsequent image has residual charges which are close to saturation. In most cases a small leftover signal in the second row is also present.

- For e2v CCDs: the first row read after an exposure that follows an exposure with saturated overscan, has residual charges which are close to saturation, and a significant signal is present in the subsequent 20–50 rows (see left-hand plot in Figure 52). The effect is slightly amplifier dependent.

These leftover electrons are not associated with what we usually call residual image or persistence. They are suspected to be associated with pockets, induced by the electric field configuration in the sensor and the field associated with saturated pixels. Investigation has revealed that only the first exposure taken after an image with saturated overscan is impacted. Our standard clear is not able to flush away those charges, while a standard readout of $\gtrsim 2000$ rows does remove them. There is a chance that a change of the electric field (e.g., a change in the clocking scheme defined in the sequencer files) can remove the pockets, and free the charges, allowing them to be cleared.

The location of these uncleared electrons in the first row of the CCDs indicates that the pockets are in the interface between the image area and the serial register. For this reason we investigated changes in the electric field configuration of the serial register during the clear, to avoid generating pockets at the image-serial register interface.

To address this clear issue, we focused on updating the serial register field as the rows are moved into it. The constraint is that the changes introduced should not significantly increase the clear execution time. It should be noted that in 2021 we tried a sequencer called “Deep Clear” sequencerV23_DC as a first attempt to address the clear issue; it added one full row flush on top of the existing one at the end of the clear. This sequencer did improve the clear, but did not fully fix the issue (see Table 4).

Table 4: Clear methods used so far

Clear Type	Duration (ms)	e2v after Saturated Flat	ITL after Saturated Flat	R01_S10 ITL "unique"
Default Clear 1 clear (seq. V29)	65.5	First row saturated signal up to row 50	First row saturated signal up to 2nd row	First 500 rows saturated for 4 amps, 13 amps with signals
Multi Clear 3 clears (seq. V29)	196.5	No residual electrons	No residual electrons	First 150 rows saturated for 2 amps, 5 amps with signals
Multi Clear 5 clears (seq. V29)	327.4	No residual electrons	No residual electrons	First 100 rows saturated for 2 amps, 2 amps impacted not measured
Deep Clear 1 clear (Seq. V23 DC)	64.69	First row saturated signal up to row <20	Tiny signal left in the first row	
No Pocket (Nop) 1 clear (seq. V29)	65.8	Signal up to row 20	No residual electrons	First 1000 rows saturated for 16 amps, 16 amps with signals
No Pocket Serial Flush (NopSf) 1 clear (seq. V29, V30)	67.0	No residual electrons	No residual electrons	First 750 rows saturated for 16 amps, 16 amps with signals

5.2.1.2 New sequencers In Run 7, we considered two new configurations on top of the default clear. The changes are in the ParallelFlush function, which moves the charges from the image area to the serial register:

- The default clear (V29): In the default clear, all serial clock voltages are kept high as the parallel clocks move charges from the image area to the serial register (sequencerV29). The charges on the serial register are expected to flow to the ground; the serial register clocks being held all high, without pixel boundaries, and with the Reset Gate of the amplifiers On. At the end of the clear, a full flush of the serial register is done.
- The No-pocket Clear (Nop): a clear where the serial register has the same configuration (S1 & S2 high, S3 low) when the parallel clock P1 moves the charges to the serial register from the image region. We kept all phases up for the rest of the time for a fast clear of the charges along the serial register (sequencerV29_Nop). The idea is that the S3 phase is not designed to be high when charges are transferred to the serial register, and is probably playing a major role in the creation of pockets.
- The No-Pocket with Serial Flush Clear (NopSf): this sequencer is close to the Nop solution, except that during the transfer of one row to the serial register, the serial phases are also manipulated to transfer two pixels along the serial register. The changes in electric field at the image-serial register interface are then even more representative of what a standard read produces, and should further prevent the creation of pockets. (sequencerV29_NopSf).

5.2.1.3 Results for standard e2v and ITL CCDs In Figures 51 and 52, we present for three types of sequencer (from left to right: V29, Nop, and NopSf), a zoom on the first rows of an ITL

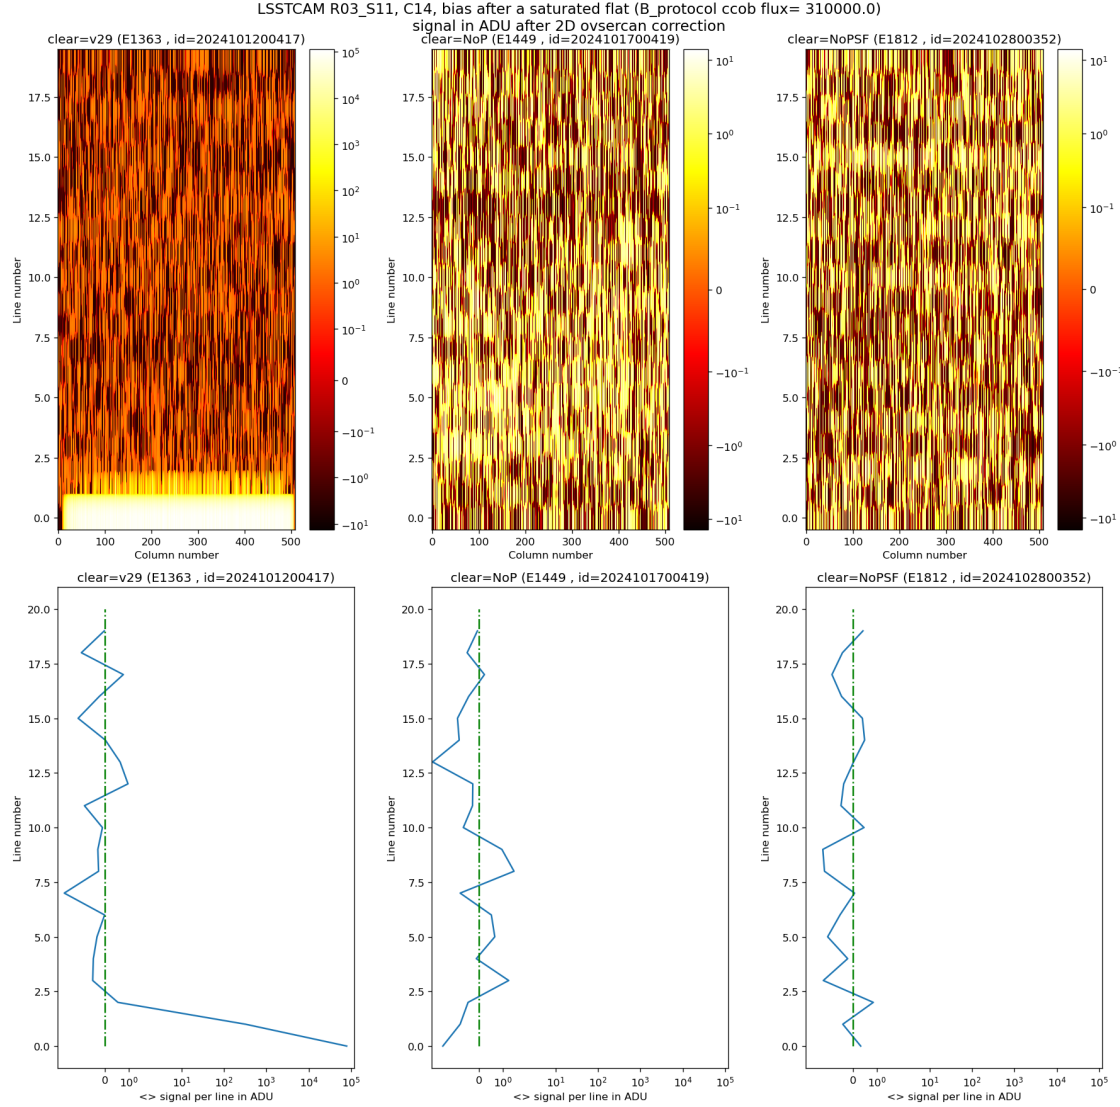


Figure 51: Same as Figure 52 but for an ITL sensor (R03_S11).

or e2v amplifier (for ITL R03_S11_C14 and for e2v R12_S20_C10 shown as a 2D row-columns image (top plots) or as the mean signal per row for the first row read of an amplifier (bottom plots)).

As seen in the left-hand panel of Figure 52 for an e2v CCD, a bias taken just after a saturated flat will show a residual signal in the first lines read when using the default clear (left images, clear= V29): the first row has an almost saturated signal (~ 100 kADU here), and a significant signal is seen up to row ~ 50 . In practice, depending on the amplifier, signal can be seen up to rows 20–50. When using the Nop clear (central plots), we can already see a strong reduction

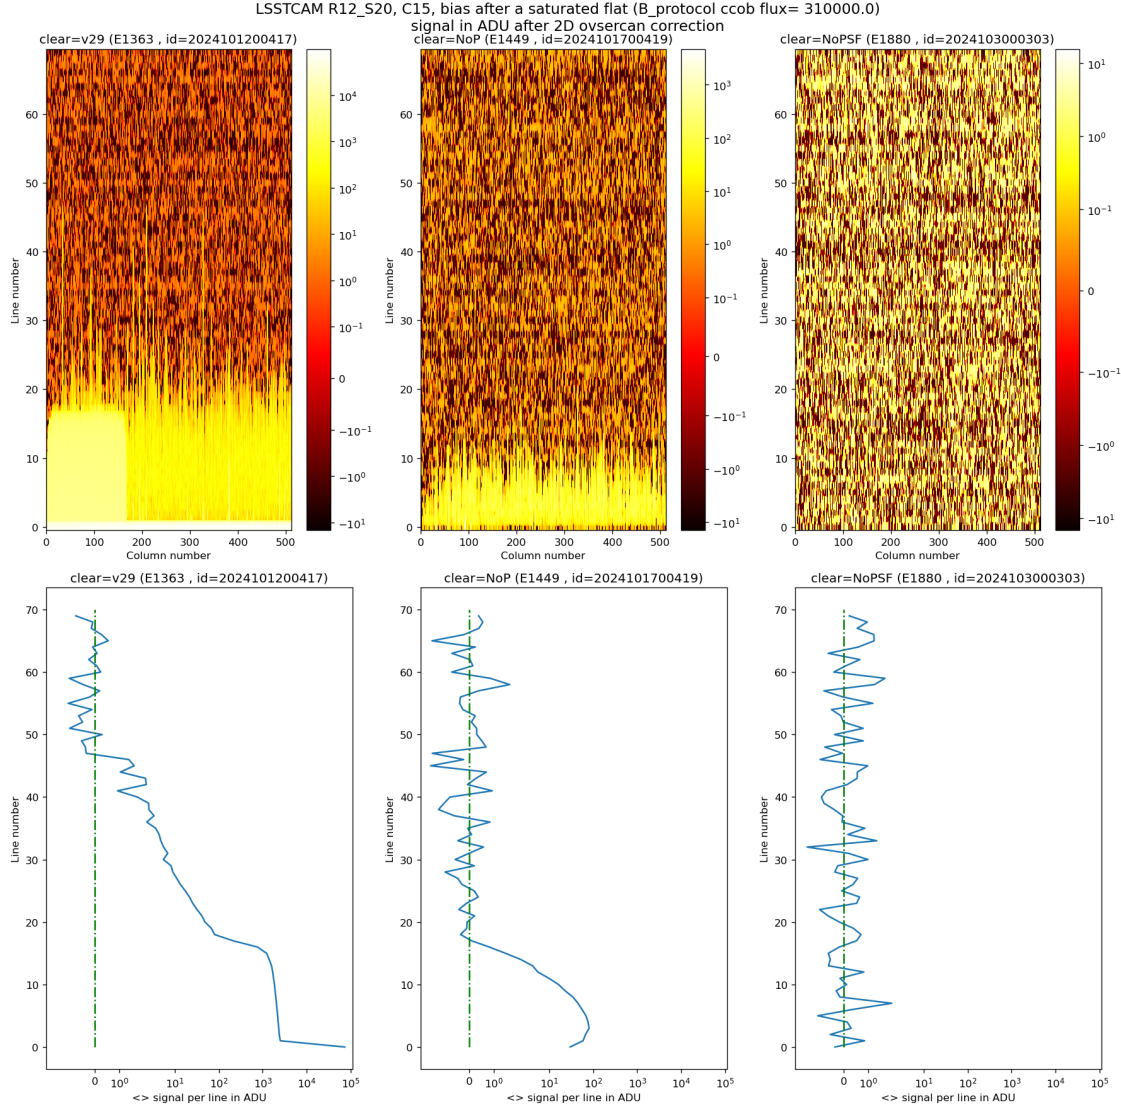


Figure 52: Impact of the three types of clear on a bias taken after a saturated flat for an e2v sensor (R12_S20). The three panels on top show the interface region between the imaging section and the serial register. The aspect ratio is not 1 for presentation purpose; the bottom three plots are the averaged column profiles.

of the uncleared charges in the first acquired bias after a saturated flat. Still a small residual signal is visible in the first ~ 20 rows. The NopSf clear (right plots) fully clears the saturated flat, and no uncleared charges are observed in the following bias.

As seen in the left-hand panel of Figure 51 for an ITL CCD, a bias taken just after a saturated flat will show a residual signal in the first rows read when using the default clear (left images,

clear=v29): the first row has an almost saturated signal (~ 100 kADU here), and a significant signal is seen in the following row. Both Nop clear (central plots) and NopSf clear (right plots) fully clear the saturated flat, and no uncleared charges are observed in the following bias image.




Figure 53: Impact of the various types of clear on ITL R01_S10 after a saturated flat (bias after a saturated flat), from left to right: 1 standard clear, 3 standard clears, 5 standard clears, 1 Nop clear, 1 NopSf clear.

5.2.1.4 An exceptional case: ITL R01_S10 One ITL sensor, R01_S10, presents a specific behavior that is not understood:

- It has a quite low full well (2/3 of nominal).
- The 3 CCDs of this REB (REB1) have a gain 20% lower than all other ITL CCDs.
- The images taken after a large saturation, as seen in Figure 53, show a large amount of uncleared charge (with the standard clear: 4 amplifiers retain ~ 500 rows of saturated signal!).

It appears that setting S3 low during the clear as done in Nop and NopSf, is even worse than a standard clear. This is strange, as a full frame read, which does this too, manages to clear a saturated image. We notice that NopSf is $\sim 50\%$ better than Nop, but still worse than the standard clear, in particular for the 12 amplifiers that are almost correct with the standard clear.

At this time we do not have a correct way to clear this sensor once the CCD heavily saturates by uniform illumination, but it is not known yet if a saturated star in this sensor, leaving signal in the parallel overscan, will present the same clear issue.

5.2.1.5 Conclusion on clears For e2v sensors, Run 7 finds the NopSf clear fully clears the leftover electrons at the interface between the image and the serial register. The NopSf clear grants that the first 50 rows of e2v CCDs that had leftover electrons from the previous exposure are now free of such contamination. NopSf will be the default clear method.

For the ITL sensors, an improvement is still needed even if Nop or NopSf overcome the clear issue because there is the significant exception of R01_S10 that prevented the usage of those sequencers for ITL devices for Run 7. Note that aside from R01_S10 the numbers of lines potentially “not cleared” in ITL devices after saturated images are small (2 first rows), and they correspond to a CCD area that is difficult to use anyway (sensor edges with low efficiency). So at this stage the original clear for ITL remains our default for ITL, i.e., serial phase 3 always, slightly extended in time to match the NopSf e2v clear execution time, will stay the default method. Further studies to overcome the problem with R01_S10 are foreseen (e.g., investigation of using a continuous serial flush during exposure at low rate, 10^6 pixel flushes in 15 s).

5.2.2 Not toggling the RG bit during parallel transfer for e2v sensors

During parallel transfer, protecting the CCD amplifier from large signal injection associated with the parallel clock swing is commonly done by activating the Reset Gate (RG) of the CCD amplifier. Although our initial default configuration, following ITL and e2v vendors’ practice at the time, this “RG protection” was not active during the last parallel clock swing. This was because following a dedicated investigation of ITL devices two years ago we found that keeping the RG active during the full parallel transfer provided a clear improvement in the biases (Utsumi et al., 2024). By analogy drawn with this ITL study and noticing that today Teledyne e2v, in its current documentation, also recommends to keep the RG active during the full parallel transfer, such approach became an area of interest for e2v devices as well.

At the end of Run 7, two runs (*E3578* and *E3628*) were collected using, for the e2v devices, the sequencer **V30_NoRG** which activates the Reset Gate during the full parallel transfer.

Although this was a limited data set a few observations can be made:

- 1D overscan bias corrected images, as they still show 2D structure, are the best place to see a first-order 2D shape change for various sequencers. For 1D overscan bias correction (serial bias overscan correction only), for detectors with large bias residuals, a

clear improvement (~ 2 ADU = $\sim 30\%$ reduction) has been observed with the **V30_NoRG** sequencer compared to the **V30** one (Fig. 55).

- For 2D overscan corrected images, no improvement was found in the stability of the bias shape with exposure time.
- For 2D overscan bias correction, leftover bias structures are unchanged between sequencer **V30** and sequencer **V30_NoRG** (Fig. 54).

So in practice running with **V30_NoRG** sequencer reduces the amplitude of the 2D bias shape of the biases: 2D residuals are smaller after 1D overscan bias correction. This confirms that part of the 2D bias shape in e2v sensors is related to parallel clock usage during readout.

Nevertheless if the biases' 2D shape is reduced, it is still present when using the **V30_NoRG** sequencer, and the corresponding issues in the removal of this 2D shape are unchanged:

- The 2D overscan bias correction presents up to ~ 1 ADU bias residual in some amplifiers (Figure 54): unfortunately there is no improvement when using **V30_NoRG** sequencer in those 2D bias residuals.
- Also the **V30_NoRG** sequencer will not be the solution to remove the dependency of the 2D bias shape with the exposure time (Fig. 56).




Figure 54: For e2v sensors, bias residuals after a 2D overscan correction: left plot, bias in run *E3380* collected with sequencer **V30**, central plot, bias in run *E3628* collected with sequencer **V30_NoRG**, right plot the difference of the two biases: No obvious difference is observed.




Figure 55: For e2v sensors, biases residuals after 1D overscan correction: left plot, bias in run *E3380* collected with sequencer **V30**, central plot, bias in run *E3628* collected with sequencer **V30_NoRG**, right plot the difference of the two biases: ~40 amplifiers show a smaller (1–2 ADU) bias residual level in the case of the *NoRG* sequencer run.




Figure 56: For 1D overscan corrected biases, here from *E3628* collected with sequencer **V30_NoRG**, we observe that the 2D shape is still a function of the exposure length: on the left we have a bias image, in the center a 15 s dark (corrected for light leak) and on the right the difference: up to $\sim \pm 1$ ADU residuals difference are observed for a few amplifiers. Notice that this has been identified as a change in bias structure as a function of the exposure time and is unrelated to dark current that is negligible for those exposure times. This is similar to what we observed in run *E3380* collected with sequencer **V30**

5.2.3 Disabling IDLE FLUSH

IDLE_FLUSH is one of the “main” settings in the sequencer file that enables the sequencer output to run while in the IDLE state (the period between one exposure and the next). The specific implementation of IDLE_FLUSH can be selected from various functions in the sequencer file. In Run 5, we chose the ReadPixel function, which reads out a pixel. This choice was initially made to mitigate the so-called yellow corner issue, a 2D structure of elevated signal near an amplifier corner observed in bias and dark exposures for certain amplifiers on e2v CCDs (see details in Utsumi et al. (2024)).

However, it was reported that running IDLE_FLUSH exacerbates the divisadero tearing issue. Divisadero tearing appears as a signal deficiency at amplifier boundaries in e2v sensors, accompanied by increased signal in adjacent columns. Additionally, using ReadPixel as the IDLE_FLUSH function has the greatest thermal impact because it continuously operates the Analog-to-Digital Converter at its maximum rate. This results in a significant difference in power consumption, more than 50 W over all rafts, between the exposure state and the IDLE state. Consequently, the focal plane experiences a temperature variation of approximately 2 deg C between periods of image acquisition and idle periods (Fig. 57).

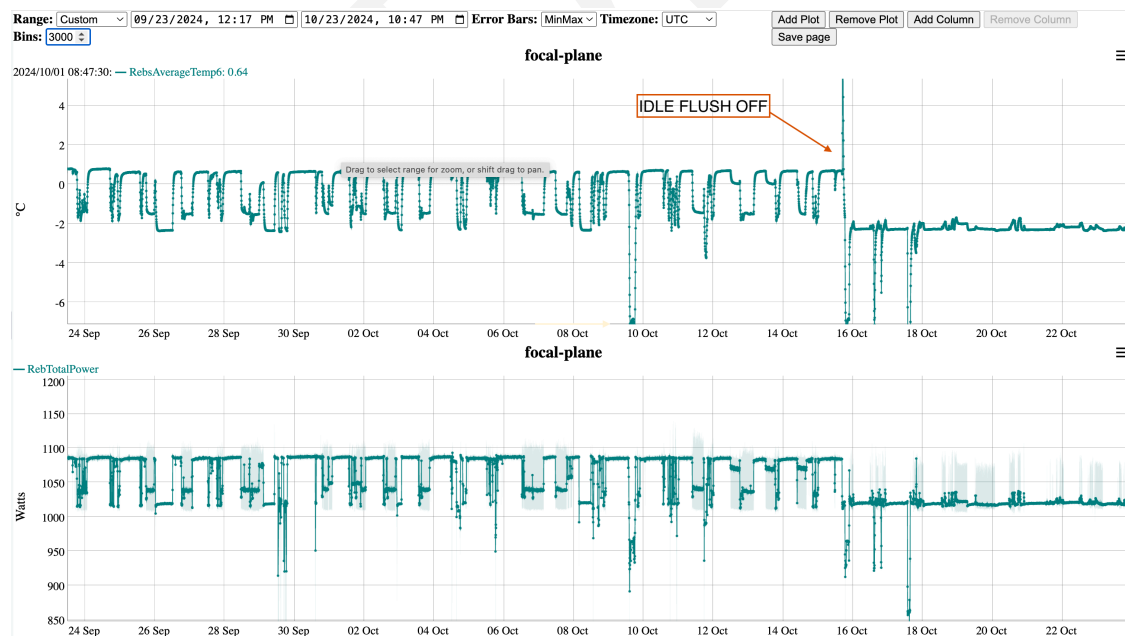


Figure 57: Impact of enabling and disabling IDLE_FLUSH on focal-plane temperature and power consumption.

This temperature variation in the focal plane can lead to changes in the REB temperature,

potentially causing gain variations or instability in the bias. Based on these considerations, we decided to disable IDLE_FLUSH. The impact of this change on bias stability is discussed in Sections 6.4 and 6.9.




Figure 58: Impact of disabling IDLE_FLUSH on divisadero tearing metric.

Figure 58 shows the impact on the divisadero tearing. Shown there are selected B protocol runs with different settings in chronological order. A few changes of settings were tested: (1) switching to narrower parallel swing voltage, (2) changing the number of clears before the exposure, (3) disabling IDLE_FLUSH. Some runs with these different settings also included some minor additional changes in settings, or a change of the sequencer file (the change from v29 to v30 is primarily incorporation of the change in the number of clears). The figure includes both ITL and e2v results. The two distinct distributions in earlier runs correspond to the differences between the two types of CCD (the upper one is e2v and the lower one is ITL). The greatest change can be seen when we switched to not running IDLE_FLUSH at E1429, which brought down the overall distribution. The two distributions became indistinguishable, which indicates that the majority of the divisadero tearing for e2v is mitigated.

E3380 was the run taken after the recovery from the shutdown due to poor performance of the Pumped Coolant System (see Sec. 8.1). The stability of the results confirms that the metric is consistent across power cycling of LSSTCam.

5.3 Summary

All e2v sensors exhibited persistence in dark images acquired after images with bright illumination. We confirmed that reducing the parallel swing voltage of the e2v CCD operation greatly reduced persistence. As penalties, we observed a full well reduction of 23% and a ~10% increase of the brighter-fatter effect in the area coefficient, essentially in an isotropic way.

Sequencer files have undergone evolution for both ITL and e2v versions. The final sequencer file from Run 6 was the v26noRG version for ITL and the regular v26 for e2v. The suffix noRG indicates that the RG bit is not toggled during parallel transfer. This modification enhances the stability of the bias structure for most ITL amplifiers.

During Run 7, several sequencer changes were implemented, as described below:

- v27 incorporated guider functionalities, including ParallelFlushG and ReadGFrame. However, the noRG change was inadvertently included. Consequently, we abandoned this version and switched to v28.
- v28 sequencer files merged v26noRG and v27.
- v29 introduced changes to speed up the guider readout.
- v30 primarily focused on e2v. We introduced NopSf for e2v CCDs. Timing of the e2v version is modified to align with the ITL version.

We also disabled IDLE_FLUSH to improve the thermal situation and the divisadero tearing.

6 Characterization & Performance Stability

6.1 Final characterization

6.1.1 Background

For a description of each quantity within this section and how it is derived from the EO test data, refer to Section 4. To compare initial and final camera metrics on Cerro Pachón, we used standard B protocol and dense red PTC data sets (Tab. 5).

Table 5: Reference runs for initial and final Run 7 comparisons

Run Type	Cerro Pachón Initial Run	Cerro Pachón Final Run
B Protocol	E1071	E1880
PTC	E749	E1881

The final operating parameters of LSSTCam for Run 7 are presented in Section 9.1.

6.1.2 Stability flat metrics

6.1.2.1 Serial CTI Serial CTI is extracted from the B protocol runs, and shows high consistency between initial and final operating parameters (Fig. 59).

The serial CTI for both sensor types is a noisy measurement, but serial CTI is not impacted by the changes to the LSSTCam operating configuration.

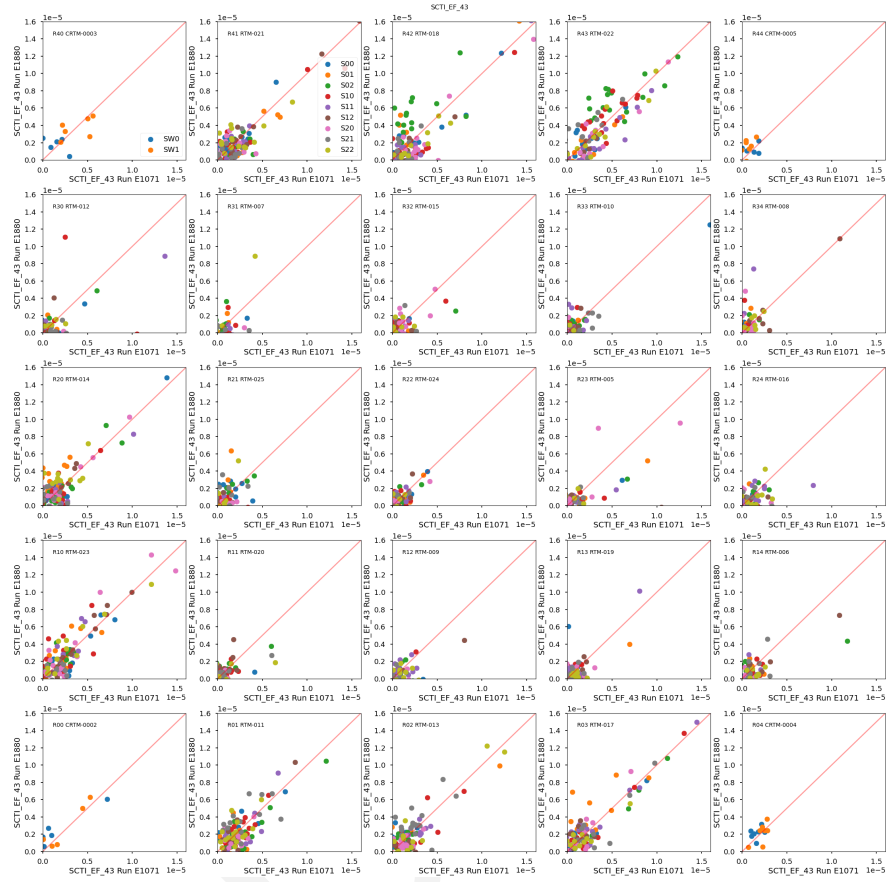


Figure 59: Comparison of serial CTI measurements for initial and final Run 7 configurations.




Figure 60: Histogram of serial CTI measurements for initial and final Run 7 configurations, by detector type. The black lines denote median measurements for the different subgroups of measurements (e2v, ITL, initial Run 7, final Run 7).

6.1.2.2 Parallel CTI Parallel CTI is extracted from the B protocol runs, and the values are consistent between initial and final Run 7 configurations.




Figure 61: Comparison of parallel CTI measurements for initial and final Run 7 configurations.

Similar to serial CTI, the parallel CTI for both sensor types is a noisy measurement, but is not impacted by the changes to the LSSTCam operating configuration (Fig. 62).




Figure 62: Histogram of parallel CTI measurements for initial and final Run 7 configurations, separated by detector type. The vertical black lines denote the median values of the different histogram populations, and are displayed in the legend.

6.1.3 Dark metrics

6.1.3.1 Dark current Dark current measurements were extracted from the B protocol runs. Across the focal plane, dark current measurements are consistent between the initial and final Run 7 runs (Fig. 63). In a subset of rafts (R13, R14, R24), a notable decrease in dark current is observed. These rafts are illuminated by the autochanger light leak, which was not mitigated until after E1880 (see Sec. 3.3.2). The physical source of the improvement in these rafts is not clear.

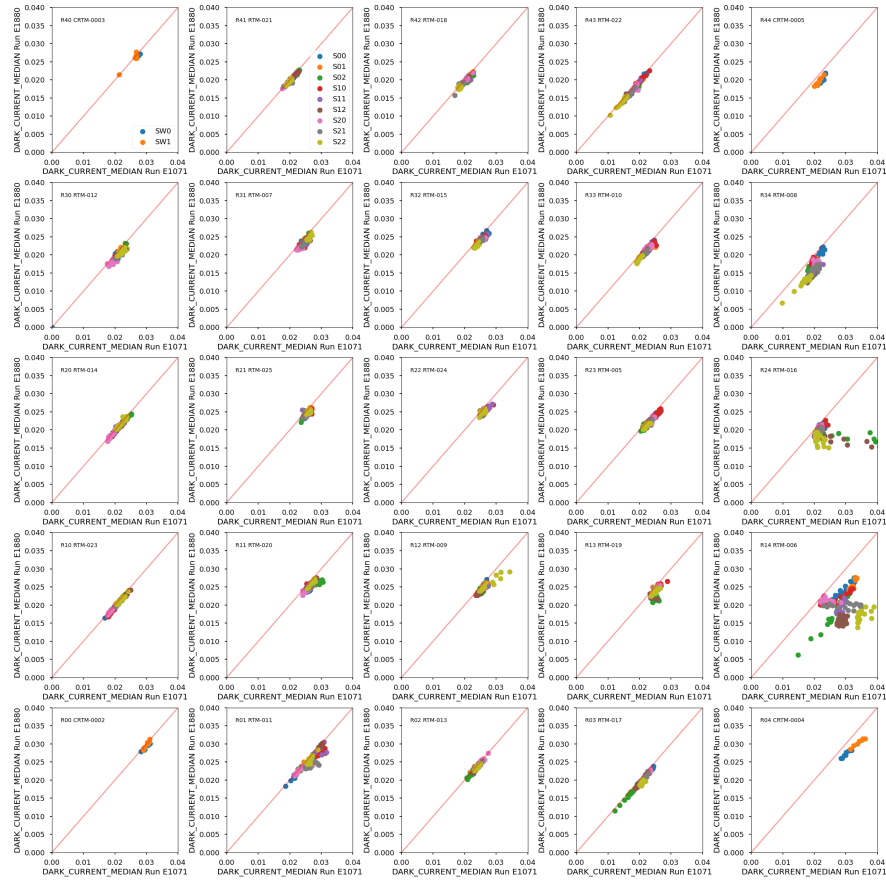


Figure 63: Comparison of dark currents for initial and final Run 7 measurements. A marked decrease in dark current is found for the rafts that were affected by the autochanger light leak during the initial run (see Sec. 3.3.2). R14 and R24 show significant changes in dark currents. The source of the improvement is not yet clear.

The reduction in dark current in the subset of rafts is indicative of light leak mitigation, and the final dark currents of those rafts are similar to the rest of the focal plane. One possible source of this improvement is improved shrouding in the region of these rafts.

6.1.3.2 Bright defects Bright defects are extracted using the B protocol runs, and the agreement between runs is extremely close (Fig. 64). No significant bright defects developed as a result of the voltage, sequencer, or idle flush condition changes.

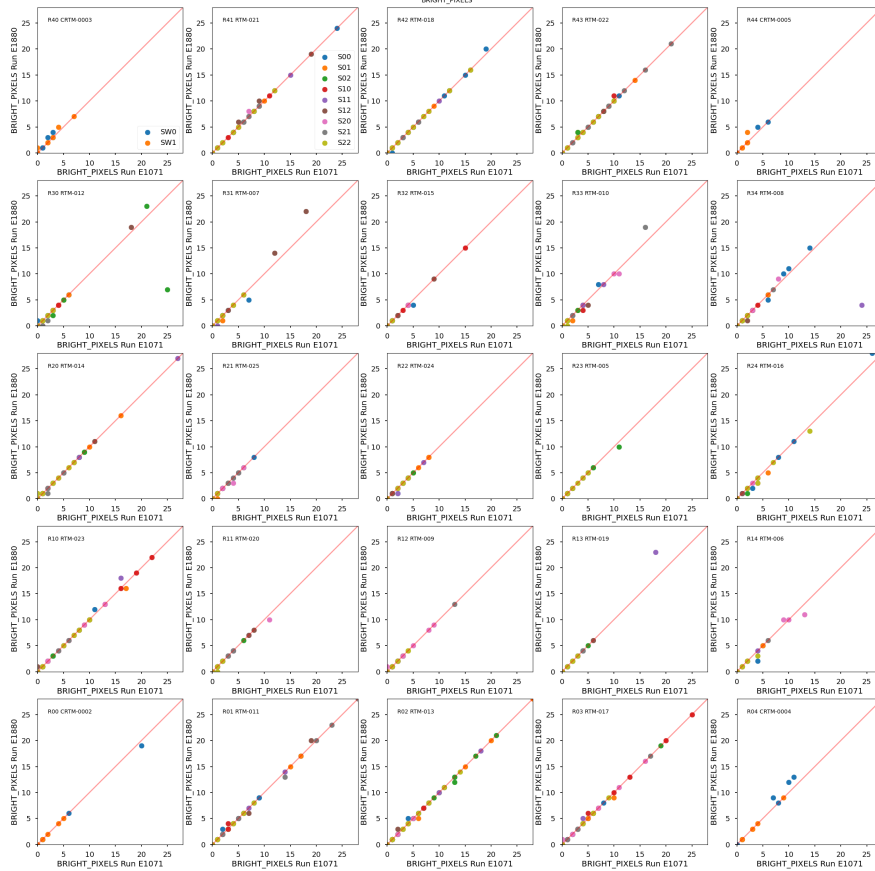


Figure 64: Comparison of bright pixels for initial and final Run 7 measurements. Each point is organized into subplots by raft. Each sensor is plotted in different colors. Measurements are made for each segment, resulting in multiple points with the same color on the same subplot.

Two segments that improved significantly were R30_S02_C17 and R34_S11_C12 shown in Figure 64. Both of these segments exhibited columns with higher response. In the final operating configuration, the spatial area of this response decreased, resulting in the improvement in the bright defect measurement.

For additional discussion about defect stability, see Section 6.3.

6.1.4 Flat pair metrics




Figure 65: Comparison of PTCs from initial and final Run 7 conditions, for R22_S11_C00.

6.1.4.1 Linearity and PTC turnoff Both linearity and PTC turnoff were evaluated from the PTC runs (Fig. 66 and Fig. 67). Due to a decrease in parallel swing voltage for e2v sensors between initial and final Run 7 runs, we anticipate a lower full-well capacity by either linearity or PTC turnoff measure (see Janesick (2001)). As described in Section 5.1.1, we observe a decrease in full-well capacity for e2v sensors. ITL sensors exhibit stable full-well capacities, despite the changes to the v30 sequencer and disablement of the idle flush (see Secs. 5.2.3 and 5.2).

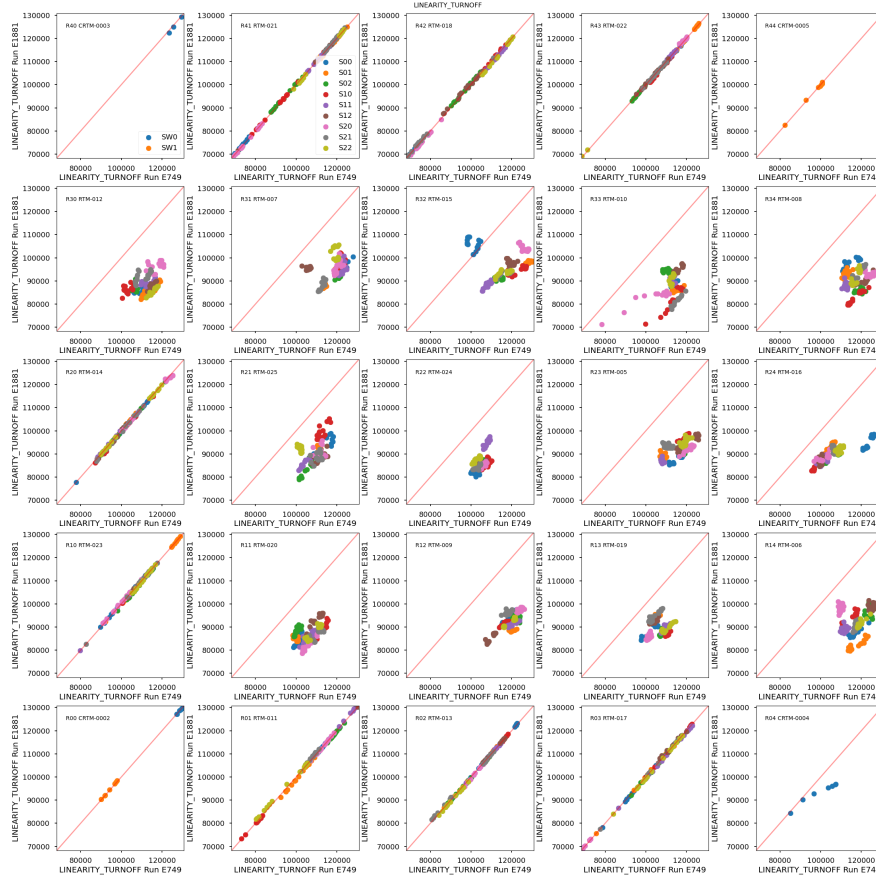


Figure 66: Comparison of linearity turnoffs between the initial and final Run 7 measurements. The e2v sensors show a notable decrease in linearity turnoff, while ITL sensors stay the same. The values reported here are in ADU, i.e., not gain corrected.

For e2v sensors we find the reduction in full-well to be significant, ~20–25% depending on the full-well metric used (see Fig. 68).

Since early in the planning for LSST, a 90,000 e⁻ full-well capacity was established as a requirement for an 8 magnitude dynamic range across all bands [Document-10512]. If linearity

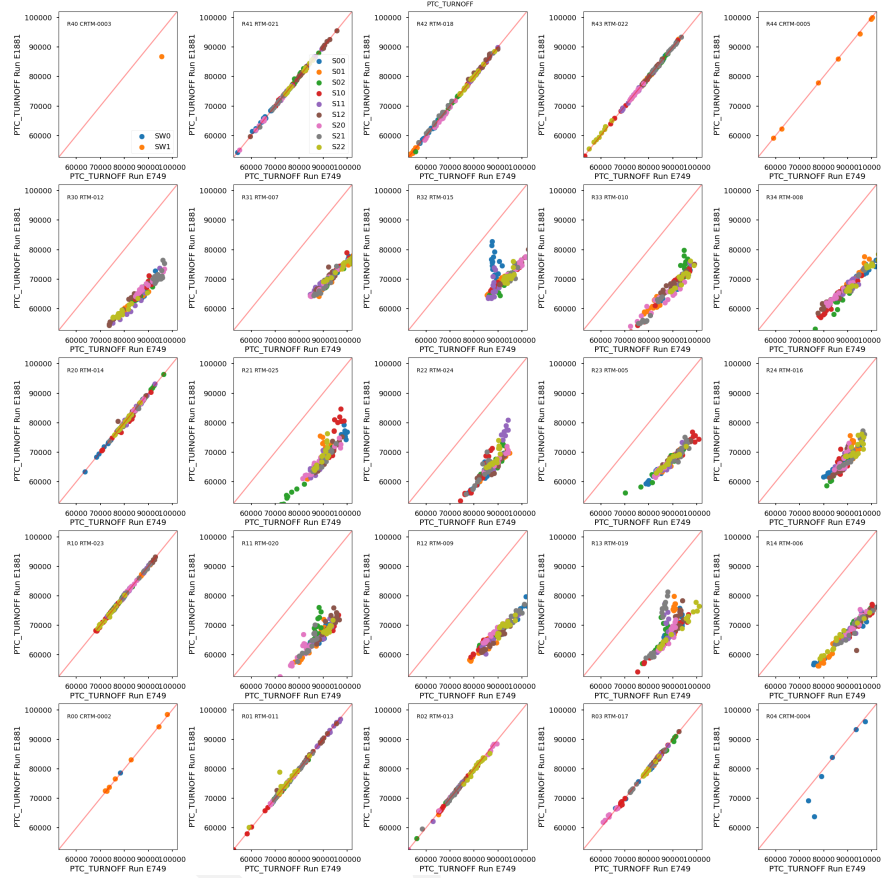


Figure 67: Comparison of PTC turnoff measurements from the initial and final Run 7 measurements. The e2v sensors show a notable decrease in PTC turnoff, while ITL sensors stay the same. The values reported here are in ADU, i.e., not gain corrected.

turnoff is the metric used to quantify full-well, all e2v amplifiers pass this system requirement. If PTC turnoff is the metric used to quantify full-well, 6.4% of e2v amplifiers do not pass this system requirement (120/1872 total amplifiers). Photometry is possible up to the linearity turnoff without a nonlinearity correction, whereas shape measurements can be made accurately up to the PTC turnoff. Shapes are not measured at high flux, near turnoff. For these reasons, we consider this requirement passed.

In PTC turnoff (see Fig. 67), several sensors (R13_S21 and R32_S00) show a feature indicative of a range of values for the final Run 7 data, but a closer distribution of values in the initial Run 7 data. The reason for this difference in distributions is not clear. Both of these sensors pass the PTC turnoff requirement.




Figure 68: Comparisons of (left) linearity turnoff and (right) PTC turnoff for e2v science sensors. The runs analyzed here are the PTC runs noted in Table 5. Note that both metrics have been gain corrected. Three peaks are observed for linearity turnoff for $dP = 9.3V$, and these resolve to one peak with a greater width for $dP = 8.0V$. A minor spatial correlation of linearity turnoff across the focal plane for the $dP = 9.3V$ configuration is the primary source of this three-peak population.

6.1.4.2 PTC Gain PTC gain was evaluated from the PTC runs (Fig. 69. PTC gain values are quite similar for the initial and final Run 7 runs, with a minor increase in gain observed in e2v sensors (Fig. 70).




Figure 69: Comparison of PTC gain amplifier measurements. The consistence between initial and final Run 7 values is quite good.

The magnitude of the gain increase for e2v sensors in the final Run 7 configuration is ~ 0.03 e-/ADU on average.




Figure 70: Comparison of PTC gains for e2v science sensors, with a moderate increase in the final run condition. The black lines denote the median value for each population of measurements, shown in the legend.

6.1.4.3 Read Noise The read noise is found to be consistent for both e2v and ITL sensors between the initial and final Run 7 conditions (Fig. 71). R01_S02 shows lower read noise in the final Run 7 configuration. R01_S02 is a known problem sensor with high read noise (see Section 6.2), so improvement in this regime is a welcome sign.

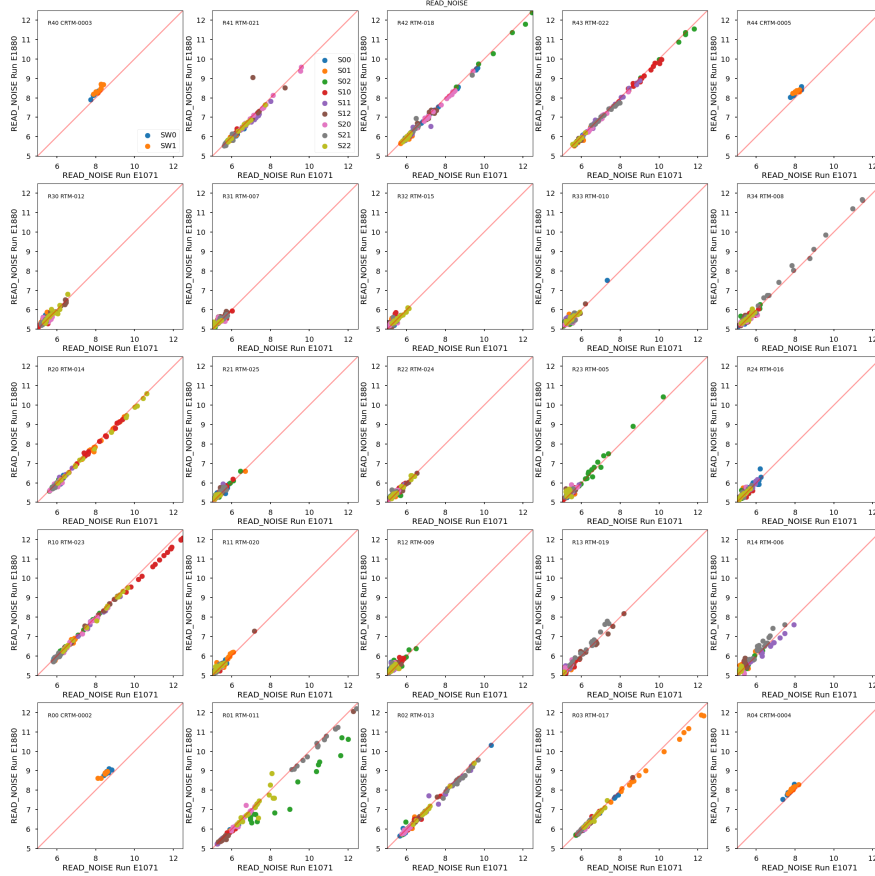


Figure 71: Comparison of read noise measurements for initial and final Run 7 conditions.

The read noise measured for both sensor types differs by $\lesssim 0.1$ e- on average between the runs (Fig. 72).

6.1.4.4 PTC Noise PTC noise is consistent between the initial and final Run 7 configurations (Fig. 73). No significant deviation in measurements is observed.

The average deviation of PTC noise measurements between initial and final runs is $\lesssim 0.1$ e- for both e2v and ITL detector types, with no significant biases emerging (Fig. 74).




Figure 72: Comparison of the read noise measurements in e2v and ITL sensors for the initial and final Run 7 runs.

6.1.4.5 Brighter-fatter a_{00} coefficient The principle area change component closely related to the brighter-fatter effect, quantified by a_{00} following the model of Astier et al. (2019), is modified in the final Run 7 operating conditions by the lower parallel swing for e2v sensors. We observe an extremely high consistency for ITL sensors (Fig. 75).

The change in the a_{00} value for e2v sensors is illustrated in Figure 76, showing a ~12% increase in the strength of the brighter-fatter effect for e2v sensors due to the lower parallel swing. For additional discussion of the brighter-fatter coefficient, see Section 5.1.3.

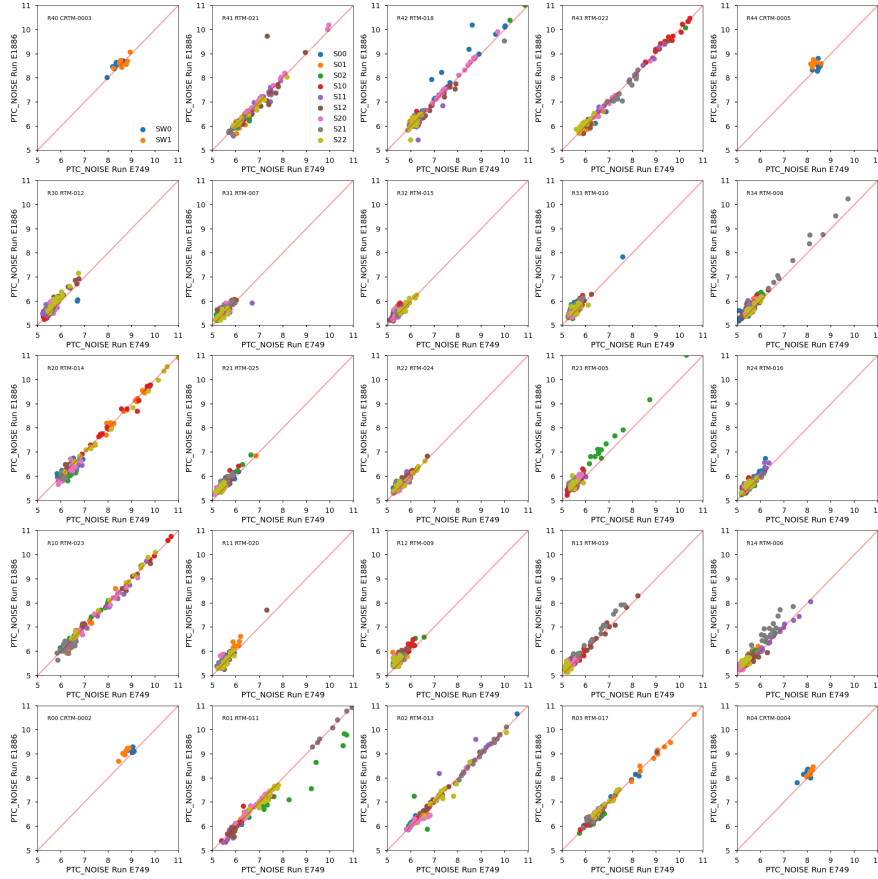


Figure 73: Comparison of PTC noise measurements for initial and final Run 7 conditions.




Figure 74: Comparison of the PTC noise measurements in e2v and ITL sensors for the initial and final Run 7 runs.

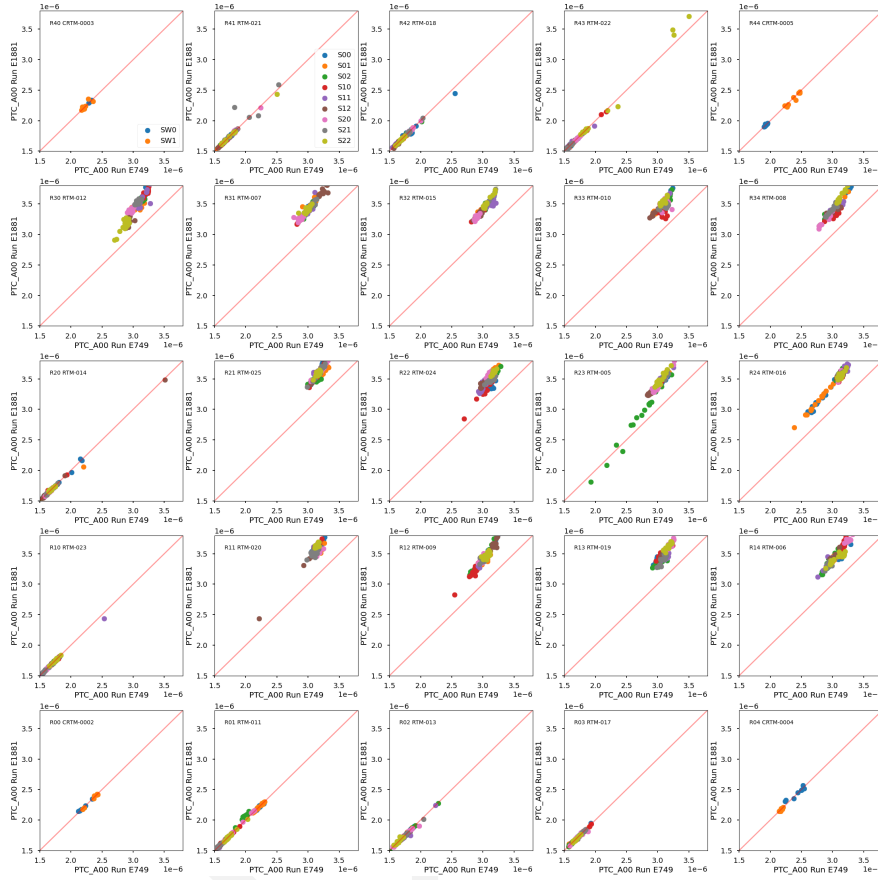


Figure 75: Comparison of per-amplifier measurements of the a_{00} parameter for initial and final Run 7 conditions.




Figure 76: Comparison of the a_{00} values in e2v sensors, showing a notable increase in the final operating conditions. As noted in the legend, the vertical black lines indicate the median values of the respective distributions.

6.1.4.6 Brighter-Fatter Correlation The strength of the brighter-fatter correlations were extracted from the PTC runs. For both x and y correlations, the results are quite consistent across initial and final Run 7 operating conditions (Figs. 77 and 78).

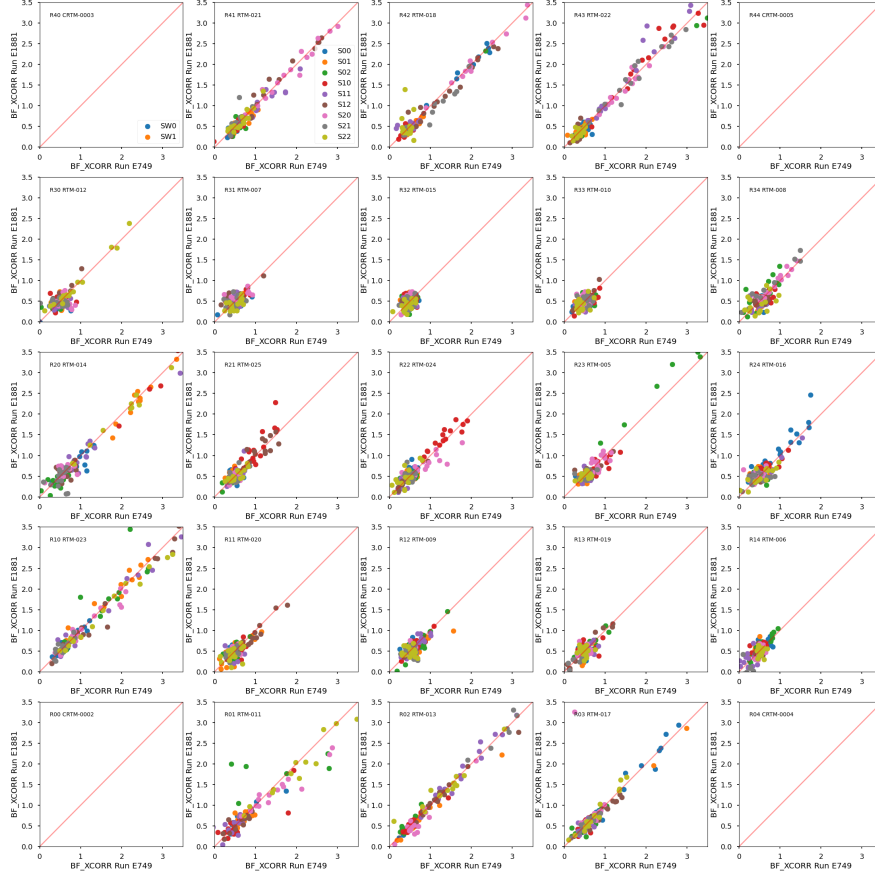


Figure 77: Comparison of per-amplifier measurements of the brighter-fatter x-correlation for initial and final Run 7 conditions. For the corner rafts, the correlations are outside of the ranges shown here.

Both correlations vary by $\pm 2.2\%$ on average, decreasing in both instances for the final Run 7 configuration (Table 6). The measurement is noisy, with all rafts showing unbiased scatter around the correlation measurement on the raft level.

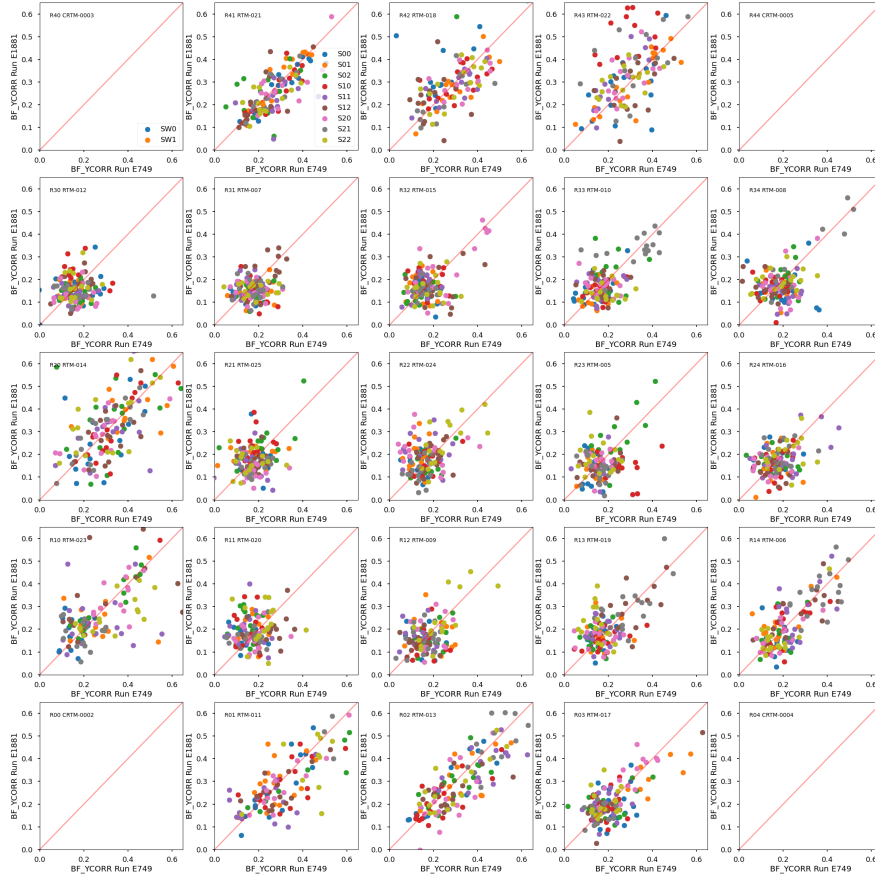


Figure 78: Comparison of per-amplifier measurements of the brighter-fatter y-correlation for initial and final Run 7 conditions. For the corner rafts, the correlations are outside of the ranges shown here.

6.1.4.7 Row-means variance The row means variance is extracted from the PTC runs, and the correlation between the initial and final operating conditions of Run 7 is very tight. The agreement is close for ITL sensors and for e2v sensors row-means variance decreases by ~1.8% in the final operating conditions (Fig. 79). Several sensors show a significant change in their final Run 7 row means variance, especially on R42 and R43 (see Fig. 79). These sensors record erroneous measurements for other PTC metrics, and warrants further study before making a determination about overall sensor performance.

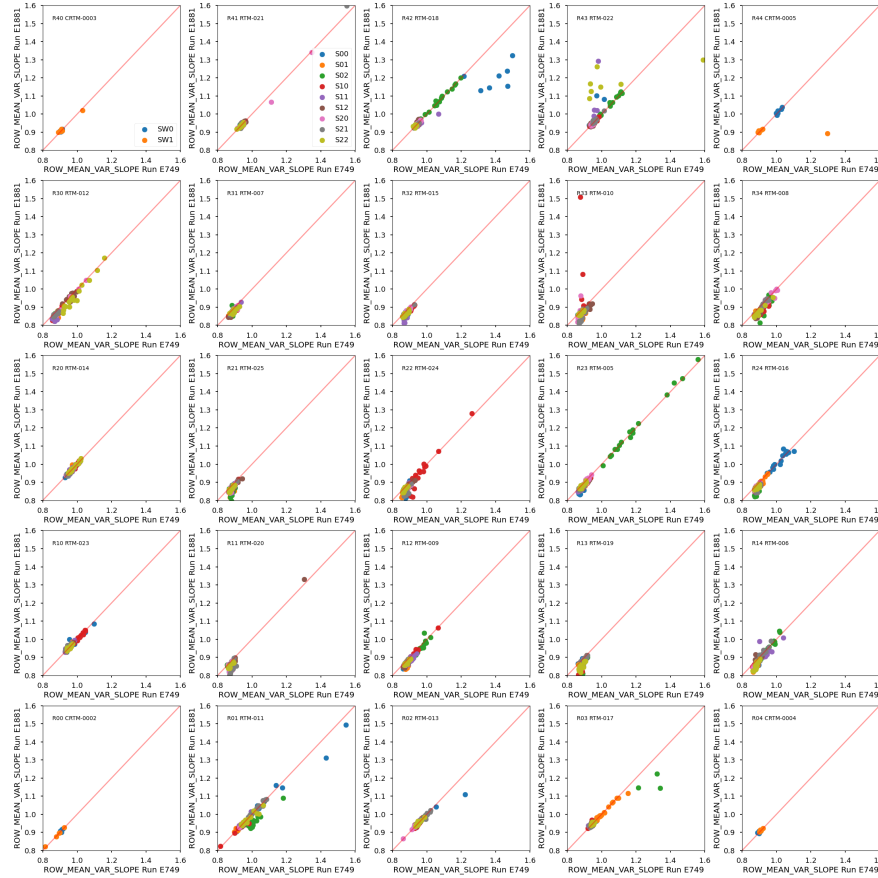


Figure 79: Comparison of per-amplifier measurements of the row-means variance slope for initial and final Run 7 conditions.

6.1.4.8 Divisadero Tearing Divisadero tearing measurements were extracted from the B protocol runs, and the results are significantly different for e2v sensors in the final operating condition (Fig. 80). The change in divisadero strength is driven by whether idle flush is enabled. Idle flush is described in detail in Section 5.2.3. The e2v sensors show a 61% decrease in the original divisadero signal under the final operating conditions. ITL sensors show virtually no change (a 0.2% increase) in the original divisadero signal under the final operating conditions.

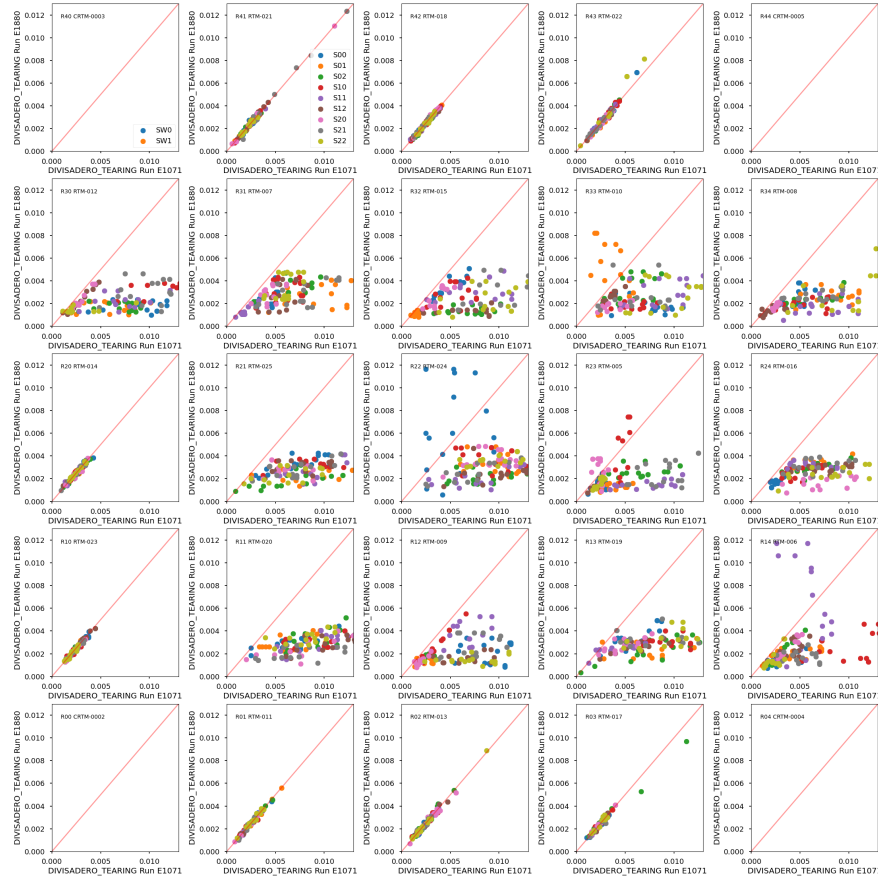


Figure 80: Comparison of per-amplifier measurements of divisadero tearing for initial and final Run 7 conditions.

6.1.4.9 Dark defects Dark defects in the LSSTCam sensors were identified using the B protocol runs. The measurement is contaminated by the picture frame effect regardless of operating conditions (see Sec. 4.4.8 for additional discussion). When applying a 9 pixel wide mask to the edges of each sensor, the picture frame signal is removed, leaving true dark defects as identified by the analysis pipeline.

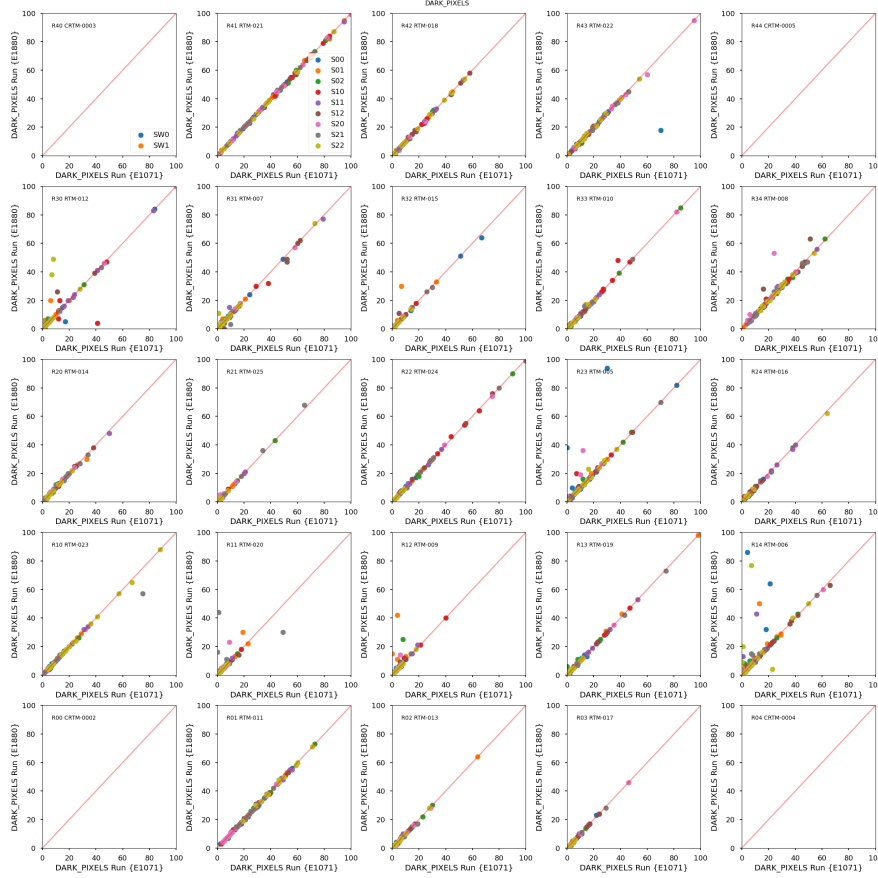


Figure 81: Comparison of per-amplifier measurements of the dark pixel for initial and final Run 7 conditions.

Dark defects are consistent between initial and final Run 7 runs (Fig. 81). Dark defects constitute a very small fraction of the pixels in the focal plane, with an average contribution of 3 pixels per e2v image segment and 8 pixels per ITL image segment. There is no global change in dark defect counts per segment, with the change of dark pixel counts per detector centered on zero for both detector types (see Fig. 82).




Figure 82: The per-amplifier distributions of measurements of dark defect pixel counts, with a 9 pixel mask applied to each sensor. Top: A histogram of the dark pixel measurements, with each count representing one amplifier. Histogram groups are separated by sensor type, and also by initial (E1071) and final (E1880) runs. Bottom: The difference in amplifier dark pixel measurements, by detector type. For both types, no significant evolution in the defect counts is evident.

6.1.5 Persistence

The primary optimization goal of Run 7 was to mitigate the persistence effect, described in Section 5.1.1. The major change in the final LSSTCam operating conditions to combat persistence is decreased parallel swing. This change is applied to the e2v sensors only, as they are the sensors that exhibit ≥ 1 ADU persistence in 15s exposure intervals with the Run 7 initial operating parameters.

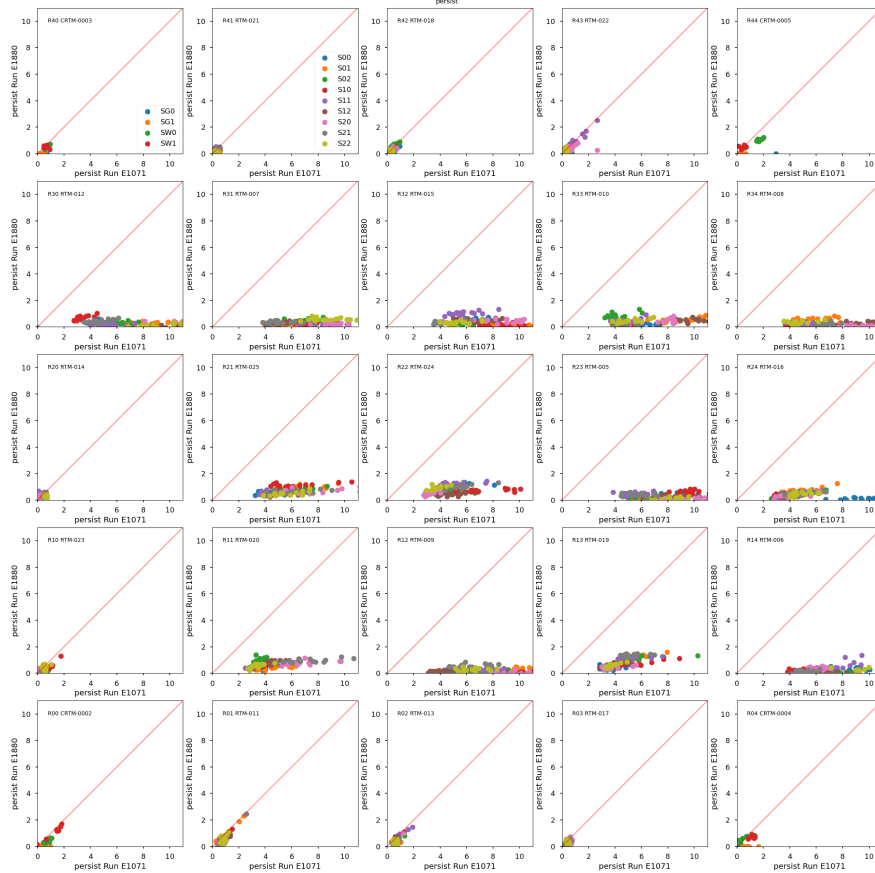


Figure 83: Comparison of per-amplifier measurements of persistence for initial and final Run 7 conditions.

The per-amplifier measurements of persistence using the metric described in Section 4.5 show a significant decrease in persistence signal in e2v sensors due to the lower parallel swing (Fig. 83), from 5.7 ADU \pm 0.4 ADU on average when measured using the red LED (corresponding to the LSST r band filter).

B protocol runs acquire persistence datasets using the red LED, which is flashed to provide

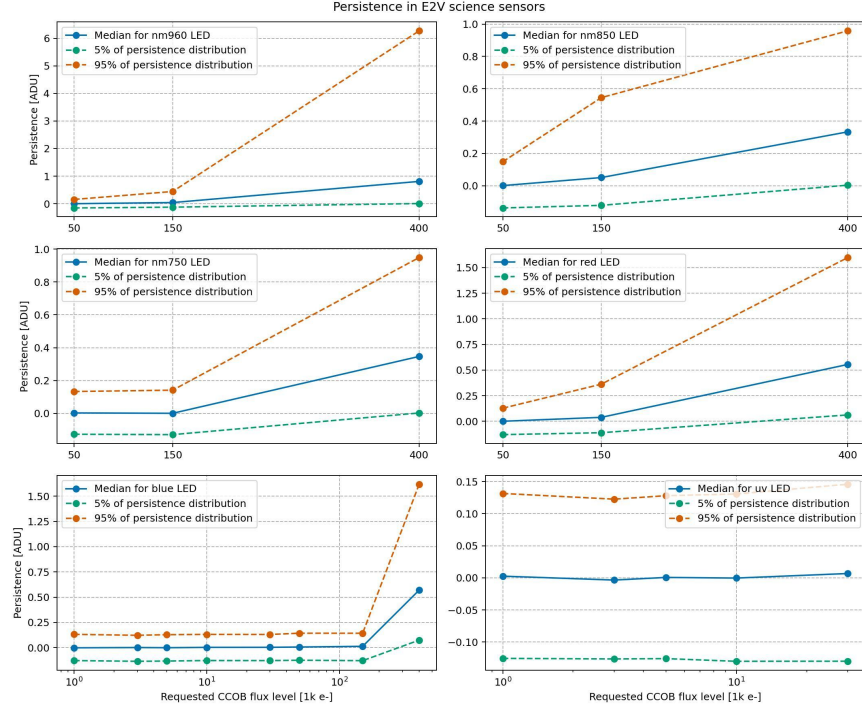


Figure 84: Demonstration of persistence mitigation in e2v sensors in the final operating condition. Persistence runs are acquired at different flat illuminations and LEDs. For e2v sensors, $\geq 95\%$ of segments show ≤ 1 ADU persistence for all exposures below linearity. The measurements of persistence below 0 ADU are due to offsets in the bias subtraction when making the frame calibration and not a feature of the persistence measurement.

equivalent of 400k e- per pixel, i.e., saturation only (a description of B protocol persistence acquisitions is provided in Sec. 4.1). Additional persistence datasets were acquired using other LEDs and other exposure levels with the CCOB wide beam projector. This was to verify that persistence was mitigated for the complete LSST photometric range. The runs used for this analysis are listed in Table 26.

We find that $\geq 95\%$ of e2v sensors exhibit a persistence signal ≤ 0.55 ADU at all flux levels below full-well capacity. The CCOB illumination varied by $\sim 10\%$ across the focal plane, and the maximum PTC turnoff for e2v amplifiers under the final operating conditions was 123k e- (see Fig. 68), within the flux levels probed in Figure 84.

6.1.6 Differences between Run 7 initial and Run 7 final measurements

Comparing the initial and final Run 7 measurements, four metrics are impacted by the optimization efforts described in Section 5.

- **Persistence:** We minimized persistence in e2v sensors, the main optimization target of Run 7, decreasing it from 5.66 ADU to 0.40 ADU on average with the red LED (LSST-r band), and maintaining sub-ADU levels across the entire LSST bandpass. Due to no change in ITL voltages and lack of an initial persistence feature, ITL sensors do not show a significant change in persistence, and remain at a sub-ADU level (0.48 ADU \pm 0.32 ADU).
- **Full well capacity:** As a direct consequence of lower parallel swing in e2v sensors, the full-well capacity of e2v sensors decreased significantly with the final operating parameters. For linearity turnoff, e2v sensors decrease from 167,796 e⁻ \pm 136,302 e⁻. PTC turnoff measurements decrease from 132,963 e⁻ \pm 102,713 e⁻. ITL sensors do not show a significant change, and remain consistent between initial and final runs.
- **Brighter-fatter strength (PTC a_{00}):** The strength of the brighter-fatter effect is also significantly impacted by the change in parallel swing for e2v sensors. The a_{00} parameter increases from $3.08 \times 10^{-6} \rightarrow 3.49 \times 10^{-6}$ for e2v sensors, a 13% increase. ITL sensors are not significantly impacted.
- **Divisadero:** The strength of divisadero tearing is impacted by idle flush. For e2v sensors, we measure a reduction in maximum divisadero signal from 0.62% \pm 0.25%, a 60% reduction in signal. ITL sensors did not exhibit a strong divisadero signal under the initial conditions, and therefore did not measure a reduction in maximum divisadero signal (0.273% \pm 0.274%). The initial strength of divisadero tearing in ITL sensors is taken as a reference size, and is therefore not minimized by the change in idle flush.

No other metrics were significantly impacted by the final operating conditions. For a complete list of the final operating conditions of LSSTCam as a result of Run 7 testing, see Section 9.1.

Parameter [unit]	Specification	e2v		ITL	
		R7 initial	R7 final	R7 initial	R7 final
Serial CTI [%]	5×10^{-4}	1.1×10^{-5}	7.3×10^{-6}	1.7×10^{-4}	1.5×10^{-4}
Parallel CTI [%]	3×10^{-4}	1.1×10^{-5}	1.1×10^{-5}	-4.8×10^{-6}	1.2×10^{-6}
Dark current [e-/pix/s]	None	0.025	0.023	0.021	0.021
Bright defects [count]	None	0	0	0	0
Linearity turnoff [e-]	>90,000 e-	168,000	136,000	178,000	178,000
PTC turnoff [e-]	>90,000 e-	133,000	103,000	129,000	129,000
PTC Gain [e- / ADU]	None	1.48	1.51	1.68	1.68
PTC noise [e-]	None	5.51	5.65	6.52	6.50
Read Noise []	<9 e-	5.32	5.40	6.26	6.21
PTC $a_{00} [\frac{1}{\text{pix}^2}]$	None	3.09×10^{-6}	3.49×10^{-6}	1.70×10^{-6}	1.70×10^{-6}
BF x-correlation	None	0.517	0.510	0.75212	0.737
BF y-correlation	None	0.171	0.167	0.287	0.284
Row-means variance	None	0.884	0.868	0.947	0.946
Dark defects [count]	<2%	3	3	7	7
Divisadero tearing maximum [%]	None	0.626	0.246	0.274	0.274
Persistence [ADU]	None	5.67	0.40	0.48	0.33

Table 6: Comparison of median parameter values on each amplifier between Run 7 initial and final measurements, separated by detector type.

6.2 List of problematic amplifiers

We classify amplifier sections as problematic if they produce effectively no signal (*dead*) when subject to incident light, or if the read noise level is above $18e^-$ (*high-noise*). Dead amplifiers are found with either read noise levels below $4e^-$ which indicates that no signal is reaching the ADC, or from anomalous PTC gain values, outside the range $1.2\text{--}2.0e^-/\text{ADU}$ (or $0.8\text{--}1.8e^-/\text{ADU}$ for BOT data, i.e., Run 5 and earlier, and single-raft testing).

A list of problematic amplifiers on Science Rafts was produced from both single-raft testing as well as a selection of runs from the BOT data-taking period during Runs 3–5 (Table 7). As the table indicates, two amplifiers (R01_S01_C00 and R10_S00_C00) transitioned from dead to working during the course of the BOT testing, and another channel (R03_S11_C00) was dead in single-raft testing, then began working during BOT testing but was dead at the end of BOT testing. At the end of the BOT testing, only (R03_S11_C00 and R30_S00_C10) were classified as dead. Furthermore, of the six channels that were flagged as high-noise during single-raft or BOT testing, only one (R41_S21_C02) remained as high-noise at the end of BOT testing.

Channel	Problem	Single Raft testing	Run 12433 2020/07/07	Run 12610 2020/10/8	Run 12795 2020/11/13	Run 12845 2021/1/27	Run 13016 2021/11/7	Run 13101 2021/11/25	Run 13137 2021/12/3
R01_S01_C00	Dead Amp	Dead	Dead	OK	OK	OK	OK	OK	OK
R01_S02_C07	HiNoise	OK	27e	22e	20e	21e	15e	14e	14e
R01_S11_C00	HiNoise	OK	24e	OK	OK	12e	OK	OK	OK
R03_S11_C00	Dead Amp	OK	NA	OK	OK	Dead	Dead	Dead	Dead
R10_S00_C00	Dead Amp	Dead	NA	OK	OK	OK	OK	OK	OK
R30_S00_C10	Dead Amp	Dead	Dead	Dead	Dead	Dead	Dead	Dead	Dead
R41_S11_C14	HiNoise	OK	NA	36e	OK	OK	OK	OK	OK
R41_S21_C02	HiNoise	OK	NA	OK	108e	96e	85e	110e	115e
R43_S02_C03	HiNoise	18e	NA	18e	18e	18e	17e	18e	17e
R43_S20_C14	HiNoise	OK	NA	OK	OK	69e	145e	OK	OK

Table 7: Table of dead and high-noise Science Raft amplifiers during BOT and single raft testing. The column “Problem” is the status *at some point during this time*. The columns to the right indicate the status of these segments during single raft testing prior to integration in the BOT and then throughout BOT-era data taking from Run 3 (July 2020) to the end of Run 5 (December 2021). For high-noise amplifiers the measured read noise is listed for levels above $12e^-$.

Next, we list problematic amplifiers detected in full Camera EO testing during Runs 6a, 6b and 7. We filter for potentially problematic amplifiers with the same cuts as above a) read noise less than $4e^-$, b) read noise greater than $18e^-$, or c) PTC gain outside the range $1.2\text{--}2.0e^-/\text{ADU}$ in a number of B sequence runs (13391, 13557, E1110, E1363, E1880, E2233, E3380) and PTC runs (13412, 13591, E1113, E1364, E1881, E2237, E3577). Note that one amplifier flagged in the BOT EO period (R10_S00_C00) is not flagged here, while there is one new amplifier (R03_S01_C05) which had never previously been flagged as problematic. For further study, the PTC and linearity plots for these eight amplifiers are shown in PTC runs in Figures 85 and

86. The eight amplifiers flagged by this selection are listed in Table 8, with comments. Note that the amplifiers listed as *dead* come in two flavors: no signal whatsoever (e.g., R30_S00_C10) or a tiny signal roughly linear with input but reduced by $\sim 10^3$ (R01_S01_C00, R03_S11_C00). The two with tiny signals had low current measurements of around 0.5 mA when they were found to be problematic, compared to 1.0 mA when they were functional.

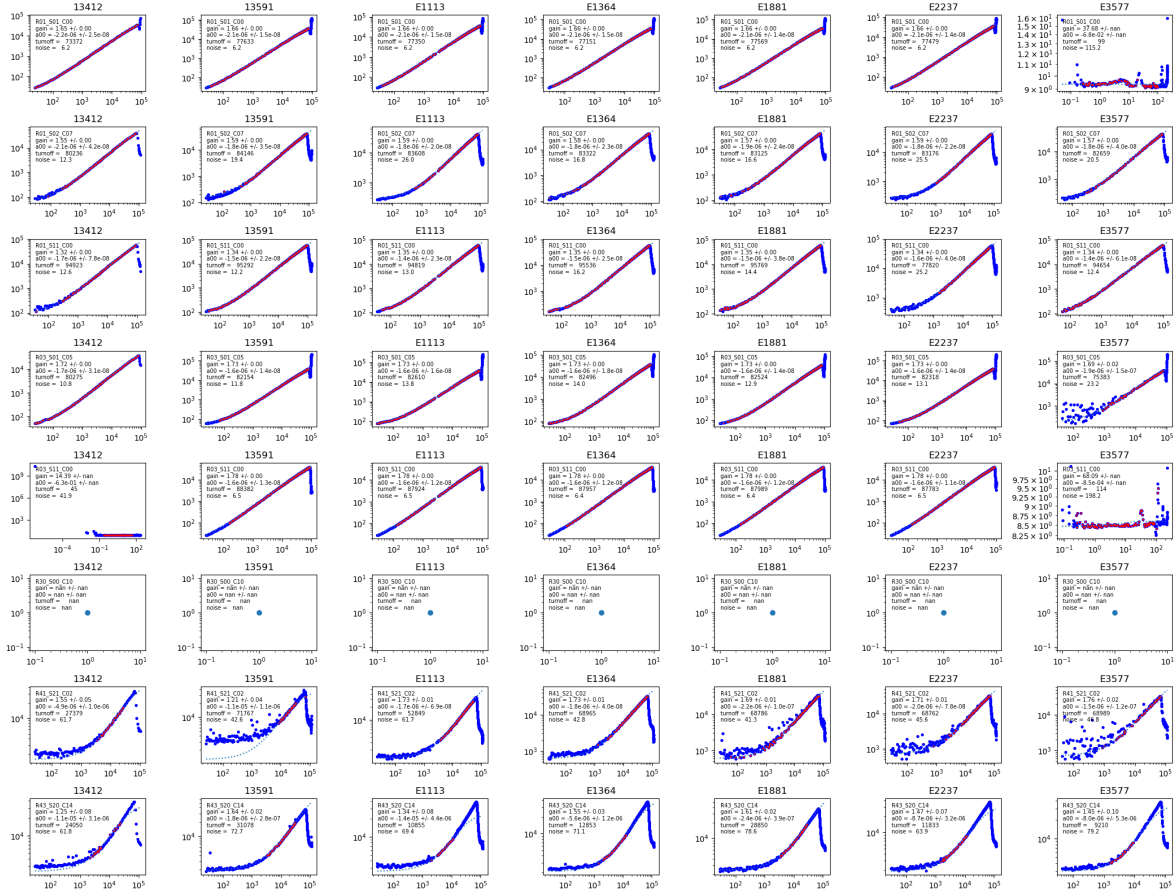


Figure 85: PTC plots for amplifiers flagged as potentially problematic, from dense PTC runs

Channel	Summary	Comments
R01_S01_C00	Sometimes Dead	Usually OK, turns dead in E3577, previously seen as dead
R01_S02_C07	OK	Noise fluctuates sometimes over $18e^-$ consistent with previous behavior
R01_S11_C00	OK	Noise fluctuates sometimes over $18e^-$ consistent with previous behavior
R03_S01_C05	Sometimes HiNoise	Previously OK, high noise in E3577 for first time, NEW bad amp
R03_S11_C00	Sometimes Dead	Usually OK, turns dead in E3577, previously seen as dead
R30_S00_C10	Dead	Always dead
R41_S21_C02	HiNoise	Always high noise
R43_S20_C14	HiNoise	Always high noise

Table 8: Table of potentially problematic Science Raft amplifiers, from Runs 6a, 6b, and 7. The Summary column gives the status at the end of Run 7. Categories are OK, Sometimes HiNoise, Sometimes Dead, HiNoise, Dead.

Finally, we list problematic Corner raft amplifiers, selected with the same filter. Three such




Figure 86: Linearity plots for amplifiers flagged as potentially problematic, from dense PTC runs

amplifiers, all in Guide sensors, have been problematic since single CCD testing. These CCDs were selected for the Guiders due to their single problematic amplifiers, rather than use a fully working Science-grade device. PTC and linearity curves for these channels, for three dense PTC runs, are shown in Figure 87, to classify these channels as either Dead or HiNoise. These three channels are listed in Table 9

Channel	Summary	Comments
R04_SG0_C11	Dead	Always dead
R40_SG1_C10	HiNoise	Always high noise
R44_SG0_C02	HiNoise	Always high noise

Table 9: Table of problematic Corner Raft amplifiers, from Runs 6a, 6b, and 7. The Summary column gives the status at the end of Run 7. Categories are as in Table 8.

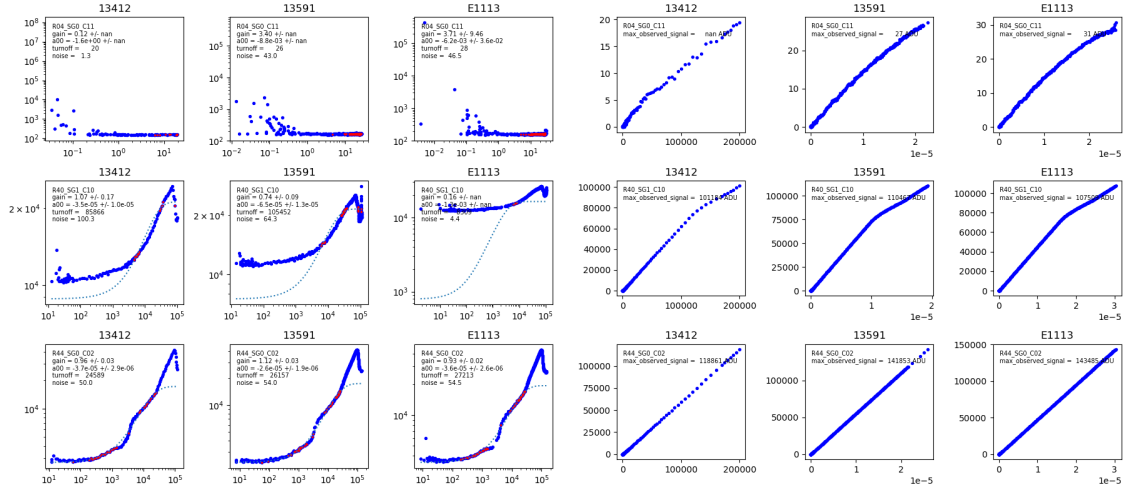


Figure 87: PTC and Linearity plots for Corner Raft amplifiers flagged as potentially problematic, from Dense PTC runs

6.3 Defect stability

EO defect masks are generated for LSSTCam images using two different protocols: one for dark defects, and one for bright defects. Dark defects masks are generated using flat calibrations, by identifying pixels with deviations greater than some threshold from the median flat value. The default dark defect deviation threshold is $\geq 20\%$ deviation from the overall median. Bright defects are identified from elevated dark current. The default bright defect condition is $\geq 5 \text{ e-/pix/s}$.

As shown in Table 6, the median number of defects in an amplifier is ≤ 10 , or $\sim 0.001\%$ of an amplifier for both e2v and ITL sensors. This contribution is extremely small relative to the useful pixels in a given amplifier. Studying any temporal dependence of defect masks on the amplifiers would be worthwhile.

To measure differences between defect masks from different runs, we can compare the resulting defect masks using binary composition. For a defect mask on an amplifier, we consider individual bright/dark defect masks from individual runs, assigning a value of 1 to each pixel that is masked. Unmasked pixels are assigned a value of 0.

Each individual defect mask can be scaled by 2^n , where $\{n : 0, 1, 2, 3, \dots\}$ denotes an index associated with a defect mask (from Run 6, early Run 7, or late Run 7). These individual defect




Figure 88: Bright defect stability over a small region of R34_S11_C06. Notably, this bright defect decreased in size as testing progressed from Run 6 → initial Run 7 → final Run 7. Left: A small region of R34_S11_C06, where the coordinates shown are in detector space for R34_S11. the colored regions denote different bright defect masks. Mask bits common to Run 6, early Run 7, and final Run 7 runs are shown in blue. Mask bits common to Run 6 and early Run 7 runs are shown in green. Mask bits common to Run 6 only are shown in red. Right: The defect count for the different mask planes over R34_S11_C06. The counts shown here are for the entire segment, not the subset of the segment shown on the left.

masks are added together to create a composite defect mask, ranging in value from $[0, 2^n - 1]$, with each bit value becoming associated with a different composition of mask planes. By comparing the defect mask bit values of different masks, it is then possible to recover on the segment level the uniqueness and stability of a defect mask by comparing the composition bit values to the mask indices. For a completely unmasked segment, the defect mask will be composed of 0's only. For a static defect mask, the composite defect mask will be composed of bit values 0 and 2^n only. This analysis is performed on dark and bright defects for Run 6, initial Run 7 and final Run 7 data (see Table 10).

Both bright and dark defects		
	Unmasked	Static mask
E2V	17.7885	57.7991
ITL	6.0764	33.1597
Bright defects only		
	Unmasked	Static mask
E2V	72.5427	93.6432
ITL	50.0000	82.3785
Dark defects only		
	Unmasked	Static mask
E2V	23.9316	60.5769
ITL	11.6319	39.5833

Table 10: Measurements of the percentages of science sensor amplifiers that meet different masking criteria, when compared across 13550 (Run 6), E1071 (Run 7 initial), and E1880 (Run 7 final). For a completely unmasked segment, the defect mask will be composed of 0's only. For a static defect mask, the composite defect mask will be composed of bit values 0 and $2^3 - 1 = 7$ only. For this tabulation we removed a 9-pixel border from each dark mask to remove the edge response from the dark defect data set. Other masking configurations were set to the LSSTCam defaults.

6.3.1 Bright defects

Bright defects are found to be stable and rare in the LSSTCam focal plane. Across detector types, 93% of e2v detectors show a static bright defect mask. For e2v sensors that had changes in bright defects, R30 dominated the subgroup with 17 amplifiers (or 11.8% of amplifiers on the raft) with a dynamic bright defect mask. Bright pixel masks of ITL sensors are more dynamic than for e2v sensors, with 18% having changes in their bright defect masks. The most dynamic bright defect masks were for R10_S02 and R10_S10, which each have nine amplifiers with a dynamic bright defect mask.

6.3.2 Dark defects

Dark defects are less stable and more abundant across the focal plane than bright defects. Across detector types, 60% of e2v detectors require a static bright defect mask. Dynamic e2v dark defect masks are distributed across the entire focal plane, with every e2v sensor registering a dynamic (i.e., time dependent) e2v dark defect mask across 13550, E1071, and E1880. Among dynamic e2v sensors, R34 dominated the subgroup, with R34_S02, R34_S12, and R34_S22 registering 15 amplifiers with dynamic dark defect masks. ITL sensors are once




Figure 89: Dark defect stability over a small region of R34_S11_C06. Notably, this dark defect developed in the final record run of Run 7. Left: A small region of R34_S11_C06, where the coordinates shown are in detector space for R34_S11. The colored regions denote different dark defect masks. Mask bits unique to final Run 7 is shown in red. Right: The defect count for the different mask planes over R34_S11_C06. The counts shown here are for the entire segment, not the subset of the segment shown on the left.

again more dynamic than e2v sensors, with 39% showing a static dark defect mask. The ITL non-static dark defect mask is present in every sensor across the focal plane, with sensors R01_S12, R01_S22, R41_S02, R41_S12, and R41_S20 measuring a dynamic dark defect mask on every amplifier.

6.4 Bias stability

We have found bias instabilities, typically above the 1 ADU level, for a number of CCDs in the focal plane, both ITL and e2v. Two main kinds of instability are observed:

1. ITL bias jumps: large variations of the column-wise structure from exposure to exposure.
2. e2v yellow corners: a residual 2D shape of the bias even after 2D-overscan correction.

These residuals depend on the acquisition sequence and the exposure time, and the

enhancement is greatest near the readout nodes (hence ‘yellow corner’).

Both issues were observed and deeply studied in Run 6 EO data. The ITL issue is believed to be random phase shifts in clocks due to the fact that REBs convert their natural clock units of 6.4ns to an artificial one of 10ns. In this conversion, there could be ambiguity in timing with respect to the natural clock. We tried to mitigate the e2v issue by optimizing the acquisition configuration in Run 7.

For the baseline acquisition configuration (see Section 9), three relevant stability runs were recorded:

1. Run E2136: 15 s darks with some very long delays throughout the run
2. Run E2236: 50 15 s darks, 50 biases recorded with 30 s delays between exposures
3. Run E2330: 15 s and 30 s darks with variable delays between exposures

To analyze these runs for bias instability, the `eo_pipe` bias stability task is used. For the ISR part, a serial (‘mean_per_row’) overscan correction is applied. The final data product of the task is the mean of the per-amplifier science image over the full set of exposures of the run. Two typical examples from Run E2136 are shown in Figure 90. In the stable case, the variations are typically at the 0.1 ADU level; in the unstable case, the variations range up to 4 ADUs.

A comparison of the results for an unstable e2v CCD (R33_S02) is shown in Figure 91 for the three runs. It must be noted that this instability can be greatly reduced at the software level with the use of a two-dimensional overscan correction (serial and parallel), as illustrated in the last panel of Figure 91, but this approach brings disadvantages whose impact needs to be estimated.

To highlight the 2D shape differences in e2v bias instability, a 2D-overscan correction is applied. A few exposures illustrating the variations of the 2D shape for the same unstable CCD R33_S02 are shown in Figures 92–94. The 2D shape of the image in amplifier C01 is different in the 3 cases.

In order to quantify the number of unstable e2v amplifiers, a stability metric d is defined from the `eo_pipe` stability task data products. More precisely, d is defined, for a given amplifier in

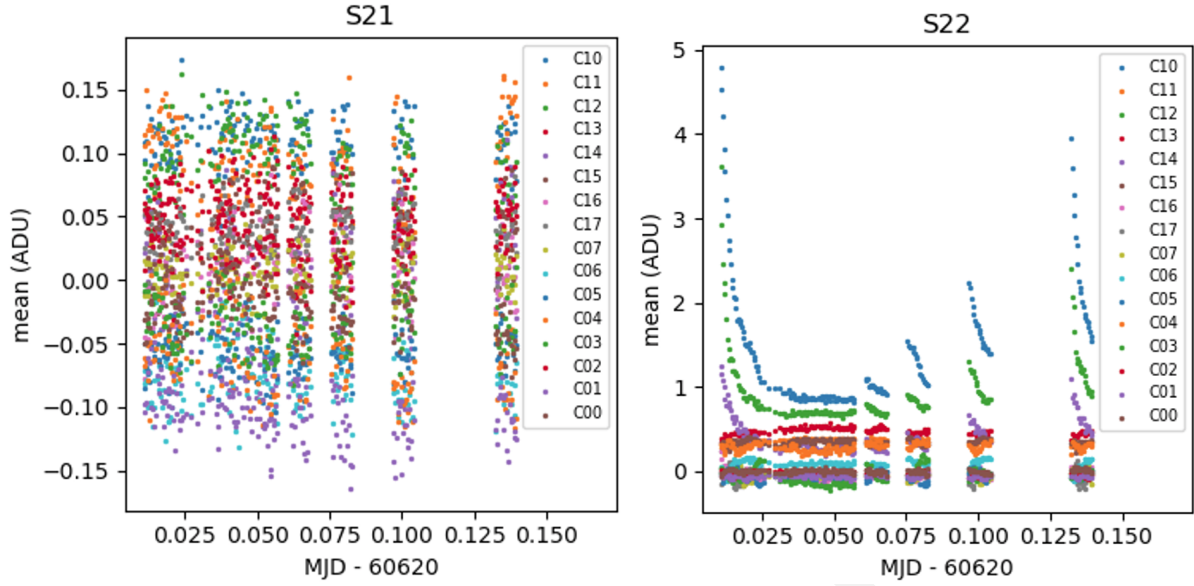


Figure 90: (left) Stable case for bias (R21_S21); (right) Unstable case (R23_S22)

a given run, as the difference between the 5th and 95th percentiles of the image mean over all the bias image acquisitions. The distribution of d for run E2136 is shown in Figure 95. Applying a threshold at 0.3 ADU (which corresponds to approximately 3 times the value for very stable amplifiers), 51 amplifiers are identified as unstable (see the corresponding mosaic in Figure 96). This corresponds to $\sim 3\%$ of the e2v amplifiers.

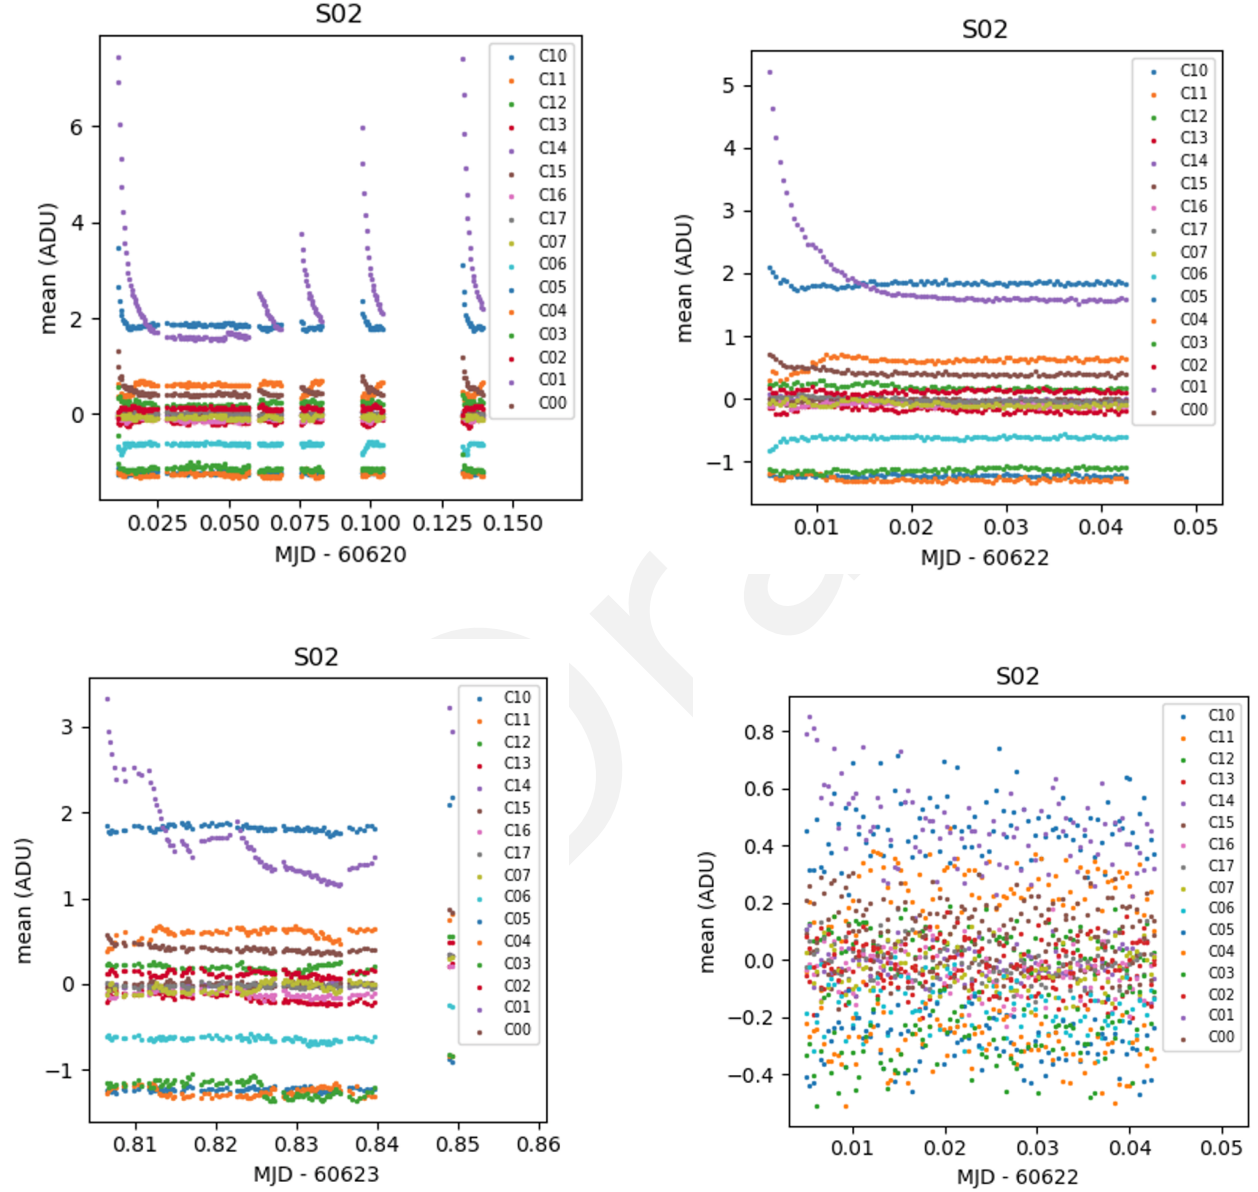


Figure 91: Bias level variations for R33_S02, an unstable e2v CCD, for three runs: (upper left) E2136, (upper right) E2236, (lower left) E2330. The segments C01 and C10 are most strongly variable in each run. Note that the range of the time axis is different in each plot. The lower right panel shows the variations for the run E2236 with the parallel overscan correction applied.




Figure 92: Bias exposure, run E1880, R33_S02




Figure 93: 15 s dark exposure, run E2136 in 'stable' conditions, R33_S02




Figure 94: 15 s dark exposure, run E2136 after a 3 min delay, R33_S02




Figure 95: Distribution of the stability metric for the e2v amplifiers in run E2136




Figure 96: Mosaic of e2v amplifiers identified as unstable (white color) in run E2136

6.5 Glow search

A long dark image can be used to identify glowing regions. Figure 97 shows an example of a 900 s dark after ISR application, with a very narrow stretch from -6 to -6 ADUs. Obvious glow can be found in 3 guider sensors of S04_SG0, S40_SG1 and S44_SG0 (see Fig. 1 to locate each sensor). The glows in the guider sensors have been identified and present since their first testing after the Raft construction. Glowing regions can be observed in R12_S22, R14_S22, R23_S11, as well, along with some other minor glowing regions elsewhere. These are the so-called yellow corner, which we discussed in Section 6.4. In R43_S21 between Segments 13 and 14, we found a very weak glow at a rate of 40–50 ADU/900 s (Fig. 98). This was also present in Run 6 data (e.g., in MC_C_20231102_000085). Otherwise, we did not find any other glowing regions in the focal plane.




Figure 97: A 900 s dark image MC_C_20240922_000432, after the ISR application.




Figure 98: A close up of R43_S21 of a 900 s dark image MC_C_20240922_000432

6.6 Illumination-corrected flat

To assess the quality of gain matching over each CCD as well as to study relative QE, we fit the CCOB-wide illumination to a smooth model and plotted mosaics of the full focal plane corrected by the illumination model. The illumination pattern is optimally fit using the 750 nm LED, since the QE is identical between the e2v and ITL CCDs at this wavelength. To capture the spatial variation of the illumination we fit the response of the focal plane using a *lambda* flat image with *physical_filter* = 'nm750' from B protocol run E2233, using cp_pipe processing with the latest ISRTask². PTC gains and linearity corrections, as determined from this B protocol run, are applied to the image and the resulting *postISRCCD* is then down-sampled into 32×32 super-pixels for ease of analysis. We take care to remove super-pixels with values in the bottom 5% or top 0.03%, values chosen by eye to remove questionable regions of the mosaic. We modeled the illumination with a two-dimensional product of Chebyshev polynomials: $I(x, y) = \sum_{i=0}^5 \sum_{j=0}^5 c_{ij} T_i(x) T_j(y)$, and fit the coefficients with the *least_squares* method in *scipy.optimize*. The 750 nm flat and the model fit are shown in Figure 99.




Figure 99: Focal Plane mosaic from a 750 nm lambda flat (left). Two-dimensional Chebyshev model illumination fit (right)

Next, we divide the mosaic images in each CCOB-wide LED by the illumination model fit to produce a relative QE mosaic, as shown in Figure 100. Visually the per-CCD gain matching looks good and we expect to refine this further by using PTC gains and non-linearity from a

²See https://s3df.slac.stanford.edu/data/rubin/lstcam/E2233/w_2025_02/

dense PTC run and by application of additional ISR pipeline tasks such as CTI, Cross-Talk and bad pixel and column masking to correct these images. We will also compare these mosaics against expectations from the single-raft measured quantum efficiency.

Draft




Figure 100: Focal plane mosaic for each of the CCOB-wide lambda flats, divided by the corresponding illumination model.

6.7 Full well measurements

One of the goals of Run 7 was to provide two initial metrics that could be used to determine what pixel value to set the SAT and SUSPECT pixel masks for on-sky data. These pixel masks flag saturated pixels (SAT) and those with levels great enough to be subject to non-linear effects (e.g., brighter-fatter effect) and need further inspection before being used (SUSPECT). For each amplifier, three values were considered: the maximum value recorded for a saturated detection, the PTC turnoff level, and the linearity turnoff (see Sec. 4.4.1 for the definitions used for the linearity and PTC turnoffs). Currently, the SUSPECT pixel mask is set by the PTC turnoff and the SAT pixel mask is set by the linearity turnoff. However, these were determined by flat illuminated images and during Run 6, we found that these values do not match what is found in saturated pinhole images (saturated pinhole values being higher (lower) than the recorded maximum values for e2v (ITL) detectors).

As such, we decided to determine the full well levels using spot images and to compare them to those found with the flat illumination setup. The spots were made using the spot projector and are the circular spots as described in Section 3.2. For these analyses, we took pairs of images ranging from exposure times of 1–100 s in 1 s intervals. To measure a spot, we applied a very loose threshold (1σ above the background) to return the footprint and when evaluating the centroid of the spot.

Once we found a spot and its centroid, we evaluated three different turnoff levels: spot photometry, spot PTC, and spot shape. For spot photometry, we used the spot centroid to place a 5 pixel^3 aperture around the spot. We then plotted the aperture flux divided by the exposure time as a function of the measured peak of the spot as shown in Figure 101. This resulted in an almost constant value until a sudden drop in the flux rate. As such, our metric for the turnoff is when the ratio first starts dropping, whenever the constant value drops below 3σ of its scatter. This turnoff is shown in Figure 101 as the black line.

For the spot PTC, we utilized the pair of images taken at each exposure level and again using the spot centroid, we selected a 10×10 pixel box around the core of the spot to use for our PTC. We then compared these 10×10 pixel cutouts between the two pairs, finding the average and the standard deviation of the sum of the two boxes. To find the turnoff level, we evaluated the

³We plotted the effect of changing the pixel aperture and found that the final turnoff value increases with aperture size. This is because this method is a proxy for measuring when the spot begins to bleed and when the bleeding exceeded the aperture size. As such, we decided to keep it small at 5 pixels.

derivative of the variance vs. spot mean and found the point that has the first large negative value (-2000 ADU/s/ADU). This curve and the turnoff are shown in Figure 101.

Lastly, we used the final full well measurement to evaluate the spot shape and its second moments. Instead of using just the centroid from the spot footprint, we utilized the corresponding bounding box to cut out the entire spot. We then ran `galsim.hsm.FindAdaptiveMom` on the image to measure the value e_1 ⁴. We then plotted this against the spot peak. Similarly to the spot photometry method, we were left with a constant value that drops off suddenly. We used an approach similar to that for the spot photometry measurements by finding the value that falls below 3σ of the constant value. The results from applying this method as well as the turnoff values are shown in Figure 101.

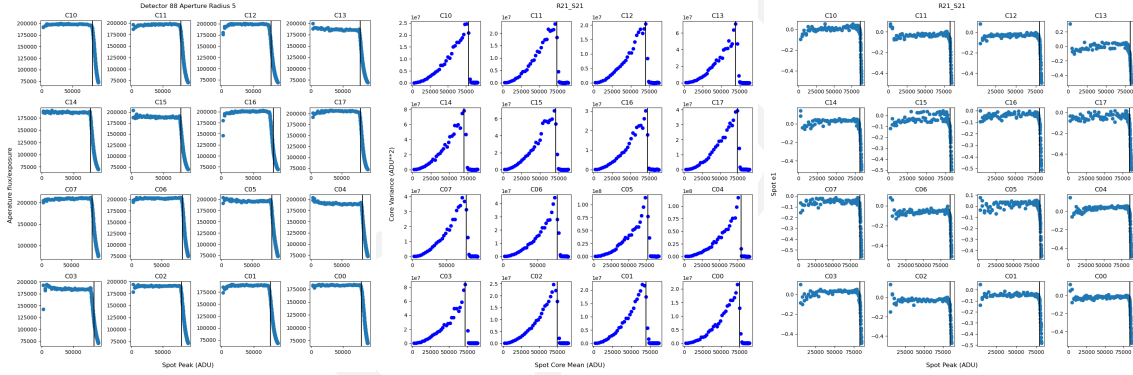


Figure 101: Example of the different spot metrics and their corresponding calculated turnoffs as the black vertical line for detector R21_S21 (88). (Left) Spot photometry, (middle) Spot PTC, and (right) Spot second moments.

We compare these results from Run 7 to the earlier results for flat images. Figure 102 shows an example comparison between the three metrics and the PTC and linearity turnoffs from the flat-illuminated data for one detector, R21_S21. This shows that all three of these metrics correspond to levels a bit greater than the linearity turnoff, providing support that the linearity turnoff is a good metric for defining the SAT pixel mask plane.

6.7.1 Spot photometry model

While R21_S21 gives good results for all three of these methods, the other detectors are not so clean, especially the ITLs. This is most likely because this detector is in the center of the focal

⁴ e_2 is also measured but we found that e_1 had more sensitive measurements, mostly likely due to it only measuring the differences in the spot shape in the X and Y directions while e_2 measures the difference in the 45 degree direction.




Figure 102: Comparison between the calculated Spot full well metrics and the EO metrics. All the Spot metrics are around the level of the linearity turnoff.

plane while the ITL sensors are arrayed in the outer parts. The data for the ITLs is messier and the same assumptions and metrics for determining the cutoffs do not work as well. To overcome this, we started using models for the data. To date, this has only been tested on the spot photometry, whose model takes the shape of

$$Flux\ rate = a - \frac{1}{bx + c}$$

where a , b , and c are constants to be solved with the model via `scipy.optimize.curvefit` and x is the peak value of the spot in ADU. We then classify the turnoff in two ways: as when the model drops by 0.1% and also when the second derivative of the model drops below a certain amount (50 ADU/s/ADU^2). Figure 103 shows an example of the data, corresponding model, and the two turnoffs.




Figure 103: Example of applying a model to the data for turnoff determination with an ITL detector (R42_S20). The data is shown in blue, the model in orange and the two turnoff determinations in red and black.

This also needs some tuning but should be less susceptible to changes across the focal plane (i.e., noisier data and different spot light levels). The work to apply this method to the whole focal plane is still on-going.

6.8 Non-linearity studies

PTC runs are meant primarily to measure variance and co-variance curves. We collect pairs of flat images, obtained using the CCOB wide-beam described in Section 3. To cover the entire dynamic range of the CCDs, we vary the length of the LED flash, the number of flashes, and the current of the LED. These data sets can be used to measure nonlinearity by comparing the CCD response to the integrated signal measured from a photodiode installed on a port of the integrating sphere that feeds a picoammeter. To avoid any shortcomings from picoammeter nonlinearity, we only compare photodiode signals of the same amplitude (illumination intensity) but different durations. We do not assume that integrated charges measured at different LED currents (and hence different photodiode currents) are on the same scale, although this turns out to be essentially true, as discussed later.

For the nonlinearity study, we use the average signal measured on each CCD channel separately, using 2D overscan subtraction and masking outlier pixels. The photodiode signal is simply bias-subtracted and time-integrated.

Technically, we model the nonlinearity using a spline function that we fit to the CCD/photodiode data pairs by minimizing:

$$Q = \sum_{ij} w_{ij}^2 \left(\frac{S(\mu_{ij}) + \mu_{ij}}{D_{ij} f_i} - 1 \right)^2 \quad (1)$$

where Q_{ij} is the CCD signal measured in exposure j at LED current i , D_{ij} is the corresponding photodiode signal, f_i is the “photodiode factor” for current i , S is the spline nonlinearity correction, and w_{ij} is some weight. We add two constraints: the average of the spline over the fitting range is zero $\langle S(\mu) \rangle = 0$, and $S(0) = 0$. If we choose equal fitting weights $w_{ij} = 1$, the residuals exhibit a scatter that varies strongly with signal level, and hence forbid meaningful outlier detection. We model the fitting weights w_{ij} using an expression determined empirically, $w_{ij} = 1/\sqrt{c^2 + v^2/\mu_{ij}}$, and the two extra parameters, c and v are also fitted by turning the least-squares expression 1 into a maximum likelihood one:

$$Q = \sum_{ij} w_{ij}^2 \left(\frac{S(\mu_{ij}) + \mu_{ij}}{D_{ij} f_i} - 1 \right)^2 - 2 \sum_{ij} \log w_{ij} \quad (2)$$

We fit the spline coefficients, the f_i factors (there are typically 3 of them), and the weight parameters c and v , for every image segment separately. We perform an iterative 5σ outlier

rejection which rejects on average $\sim 0.5\%$ of the data points (this small rate validates the modeling of weights). Figure 104 displays some results of the fits. The quality of the measured non-linearity is sufficient for our needs.

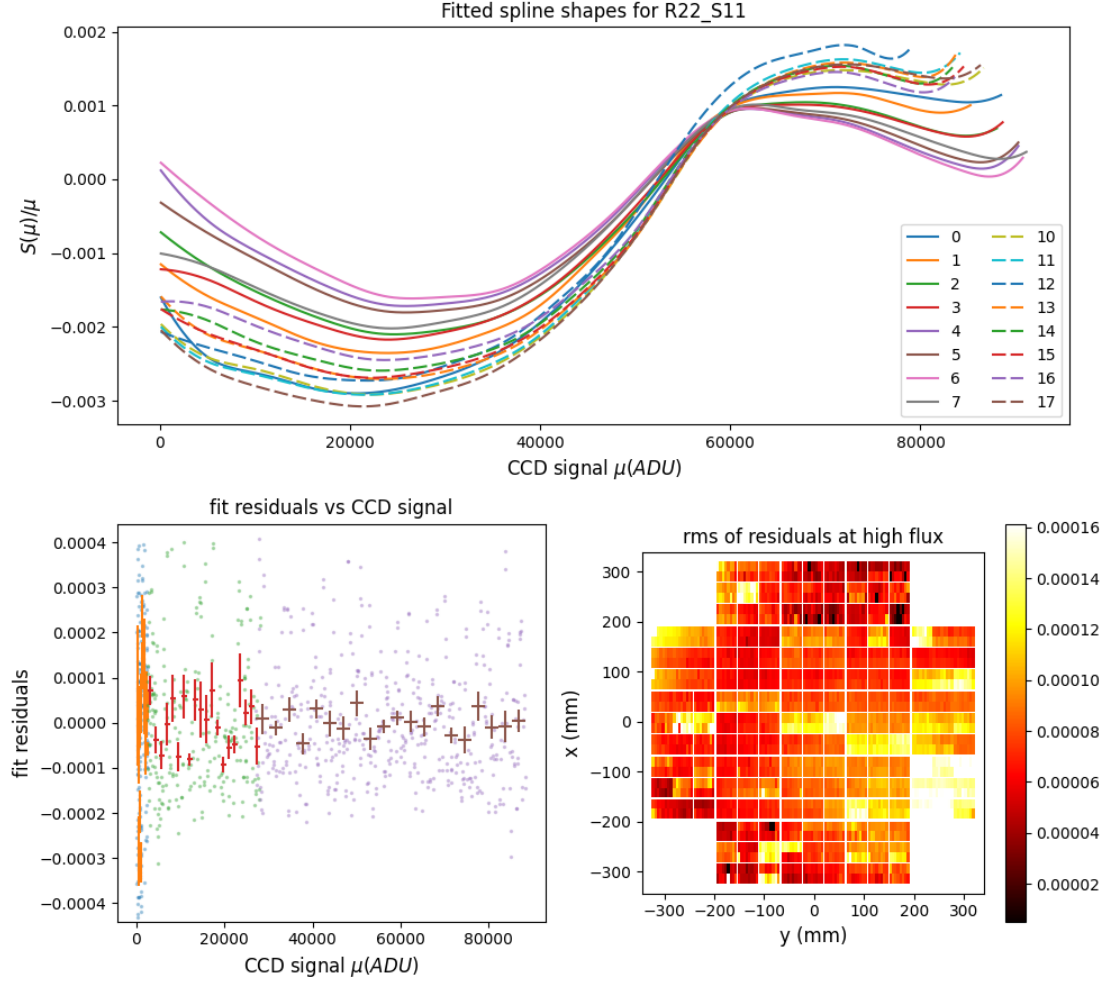


Figure 104: top: fitted nonlinearity spline (divided by the signal level) for the 16 channels of R22_S11. The main feature is due to the distortion introduced by the preamplifier. Left bottom: The fit residuals for channel 0 of the same CCD. Different colors refer to different LED current settings. Bottom right: r.m.s. of high-flux fit residuals (the c parameter of the fitted dispersion model) for all camera channels. Those are about 10^{-4} on average, and some are correlated within REBs, for an unknown reason. The quality of the obtained correction is well within goals.

6.9 Gain stability

We have used “relative gain” in this section to study gain stability over time. The relative gain is defined as the ratio of the signal observed in a CCD image segment divided by the integration of the photodiode current with respect to an arbitrary normalization. With a fixed flat illumination, the variation of the relative gain over successive exposures can be utilized to investigate the gain stability. In the past Run 6 (Utsumi et al., 2024), we used a run during which the Cold plate target temperature was kept constant, which reflects the real observing condition. We repeated this test during Run 7 and we acquired flat images at two representative flux levels with two distinct temperature conditions: either intentionally altered or maintained constant.

- E1496 (dp80, 12 hr, constant temperature, v29_Nop, nm750, 10k e-)
- E1367 (dp80, 6 hr, temperature swing of 2 degrees, v29, nm750, 50k e-)
- E756 (dp80, 6 hr, gain stability @ 50k e-), unprocessed because there was a data transfer issue
- E1362 (dp80, 6 hr, 10k e-), unprocessed because there was data transfer issue

Here we focus on E1496 (steady temperature) and E1367 (variable temperature).

Figure 105 shows the derived relative gain change for one amplifier (R11_S00_C11) over time along with other representative parameters such as the average of two REB temperatures measured closest to the ASPICs, $(TEMP6+TEMP10)/2$, CCD temp, and LED1TEMP (the temperature measured at the LED board on the CCOB projector). In Run 6, the REB temperature average was found to be a good proxy for the relative gain change. For the steady temperature acquisitions, the REB temperatures were maintained steady within 0.2 C, with a slight decreasing trend during the run probably due to the change in the thermal load and the stabilization process of the entire thermal system. During the run the gain slowly increased over time, while the CCD temperature remained steady and the LED temperature was not controlled but stayed within an 0.1 C range.

Figure 106 shows the relationship between the relative gain and the REB temperature, with color coding by the rate of temperature change, along with its fit. A reference line from the past result is overlaid with an arbitrary vertical offset. Clearly, the gain-temperature relationship is steeper than the previous result. The distribution of the data points has a more complicated




Figure 105: relative gain changes with other parameters for one amplifier R01_S00/C11 in run E1496.




Figure 106: Relative gain as a function of REB temperature (average of TEMP6 and TEMP10), with color based on temperature swing run E1496. The fitted line and the line with the slope taken from Run 6 are also displayed. The intercept of the Run 6 line is arbitrary, since the operating temperature was different.

structure than the linear relationship, while there is no obvious change in either CCD or LED temperatures.

Figures 107 and 108 show the same set of plots but for the run that has a cold plate temperature change of 2 C. The temperature was kept steady in the beginning, but the set point was changed later, and then it was brought back to the original temperature. Clearly, Figure 108 shows not only temperature dependency but also hysteresis in the gain-temperature relationship, which does not match the slope originally derived from past runs, although there are no obvious changes in the system other than the REB temperature. During this run the LED board temperature also decreased by more than 1 C over the testing period.

The reason why the relationship between gain and temperature became much more complicated is not clear. It is understandable that hysteresis was not observed in Run 6 because there was no intentional temperature change in the cold plate, which means the cold plate/REB temperature swing was minimal. However, looking at the result from E1496 where we took images at the same temperature, the relationship is much more complicated than what it looked like before. A number of possibilities can be considered to explain this: 1) there is a hidden variable that changes the gain other than the REB temperature, 2) illumi-




Figure 107: Same as Figure 105 but for run E1367.




Figure 108: Same as Figure 106 but for Run E1367.

nation from the LED is somehow changing over time, which is not correlated with the LED temperature, 3) air turbulence in the lens volume contributes to this, or 4) condensation on the lens might come into play. As hysteresis is observed, possibility 1) is definitely relevant but the particular hidden variable is not clear because it cannot explain the gain changes when the temperature is constant. Possibility 2) is unlikely given the fact that the complicated relationship was observed as well in Run 6. Possibility 3) could play some role since Banovetz et al. (2024) discovered illumination changes due to turbulence of the air in the lens volume. However, it is not clear if any kind of long-term trend over 6 hours can be explained by this. For possibility 4), we made a visual check for condensation (not during the run) and we did not find anything obvious.

Gain changes can be categorized in two ways: global or local to an amplifier. A global coherent change can be, in principle, correctable as it is degenerate with atmospheric transparency variations, which will be corrected by the calibration process. A local amplifier-by-amplifier change would be a more serious issue because the number of stars in an amplifier segment might not be sufficient for making precise photometric calibration. In order to study the local amp-by-amp gain change, Figures 109 for the constant temperature condition and 110 for the temperature swing condition show the differential gain changes with respect to the medianed relative gain for the entire focal plane.




Figure 109: Time dependence of relative gain change relative to the overall median gain, for E1496

The differential gain change with respect to the global change for the constant temperature case appears mostly stable within the level of 10^{-4} . Some of the measurements deviated from zero because of the normalization of the first measurement. R11_S12, R12_S10, R12_S22, R24_S11, R34_S20 have one amplifier that has a higher relative gain up to 5×10^{-4} but others behave stably. This could be contaminated by the yellow corner in e2v sensors (Sec. 6.4) but

this could be mitigated by discarding the first few exposures, which is probably required for other reasons such as bias instability. Further investigation is needed. Another interesting behavior is seen in R11_S2x. The three sensors had a spike of gain simultaneously. We have not figured out what happened at that time.




Figure 110: Same as Figure 109 but for E1367

Interpreting the results from E1367 with the temperature swing is complicated. Some of the amplifiers behave as well as the ones for the constant temperature case, but others show

correlation with the temperature change in the Cold plate temperature. This indicates that a relative gain change among amplifiers with respect to REB/Cold plate temperature exists. Note that the flat images for E1367 have a 5 times great flux than E1496, which reduces shot noise in the measurement. However, the conclusion still holds.

To further study gain stability, we stepped back to the raw measurements. Figure 111 shows the constant temperature case E1496. The change in the relative gain is at a level of 2×10^{-4} , which appears to be driven by the photodiode integration. Figure 112 shows the temperature swing case E1367 with a change of 5×10^{-4} , which appears to be dominated by a change in image counts. The changes in the PD integration are about the same in both plots. So from these facts, both a gain change in the Camera due to the temperature change and some illumination difference of the CCOB projector play roles here.




Figure 111: Raw measurements of image count and photodiode integration, as well as the ratio of those – the relative gain for E1496




Figure 112: Same as Figure 111 but for E1367

In summary, we find

- The gain-REB temperature relationship is not as simple as it was in Run 6.
- Global gain change could be due to artifacts or the setup, or potentially the Camera could have a complicated behavior with respect to the REB temperature. No conclusive statement can be drawn.
- Local amp-by-amp gain change is minimal 10^{-4} over 6 hours if the REB and the focal plane temperatures is maintained steady.

- Further analysis is needed for understanding the gain change in the beginning of Run 7 and the apparently random spikes of gain.

Draft

6.9.1 Non uniformity in focal plane response change

In the long stability runs most of the observed variations are common to all amplifiers over the focal plane. We identified 3 sources of common mode effects:

- variation of the intensity of the light source itself
- a drift of the photo-diode measurement
- a coherent temperature change of the REB electronics

In practice the CCOB LEDs, only sensitive to their monitored temperature, appear to be extremely stable during Run 7. In the long stability run it seems that the CCOB light source was more stable than the measured integrated photo-diode current that we used to correct it; e.g., in run E1496 a response drift of 0.3% over the run appears only when the photo-diode correction is applied to the data.

From Run 6 studies we know that the usage of IDLE_FLUSH, that warms the electronic when not acquiring data, induces a coherent cool down of the REB electronics at the start of each data-taking run. So the usage of IDLE_FLUSH can generate a coherent increase of the electronic gain at the start of the run ($\sim 6.5 \times 10^{-4}$ /deg C). Due to this effect, we had up to +0.1% (~ 2 deg change) gain increases in the worst case during Run 6.

In practice a coherent mode in signal response over the focal plane does not really matter for science; it will be absorbed in the common gray term measured with stars for each exposure. In the rest of this section we will focus on the changes in response that are not uniform over the focal plane.

6.9.2 Non uniform temperature variation.

In Figures 113 and 114 we present the temperature variations measured per REB during two long runs:

- Figure 113, run E1496, IDLE_FLUSH active: we saw the common drop of the temperature at the start of the run, it's in general extremely small (-0.1 deg), except for some REB (ex R24/Reb2 has a drop of 0.6 deg at the start of the run)

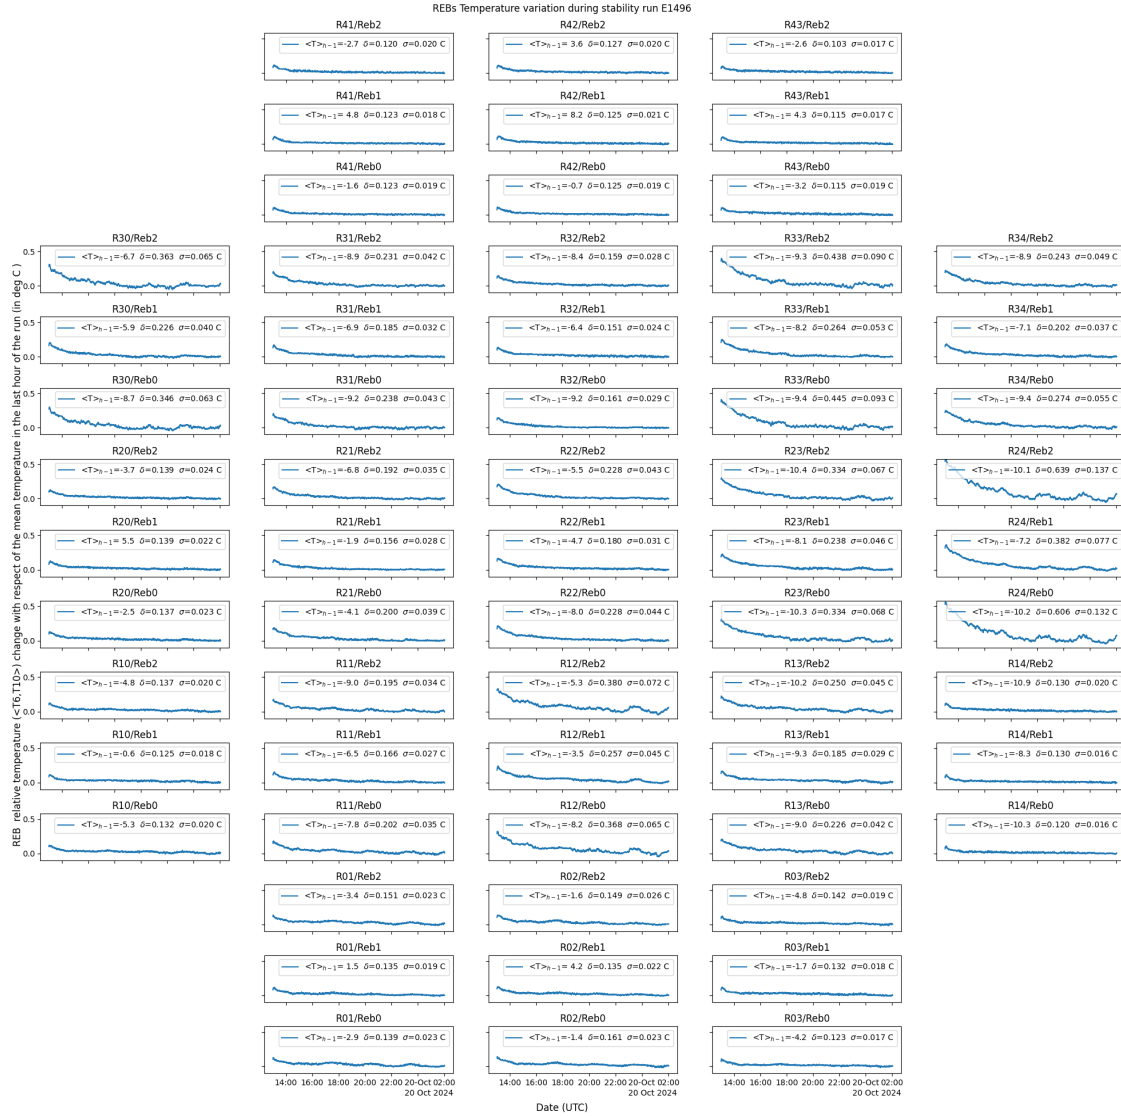


Figure 113: REBs temperature change during run E1496 (IDLE_FLUSH On). In the plot the temperature of each REB is relative to the mean temperature during the last hour of the run. For each REB the value of this mean temperature ($\langle T \rangle_{h-1}$), the peak-to-peak temperature change (δ) and the temperature dispersion over the run (σ) is given. We clearly see that not all temperature variations are the same; for example the REBs of R24 have larger variation than the ones of R14

- Figure 114, run E2008, no IDLE_FLUSH, there is for most of the REB a temperature dispersion < 0.02 deg C, except for some REBs (e.g., R24/Reb2) which have dispersions more than 3 times greater (a kind of oscillation with peak-to-peak variation of ~ 0.3 deg)

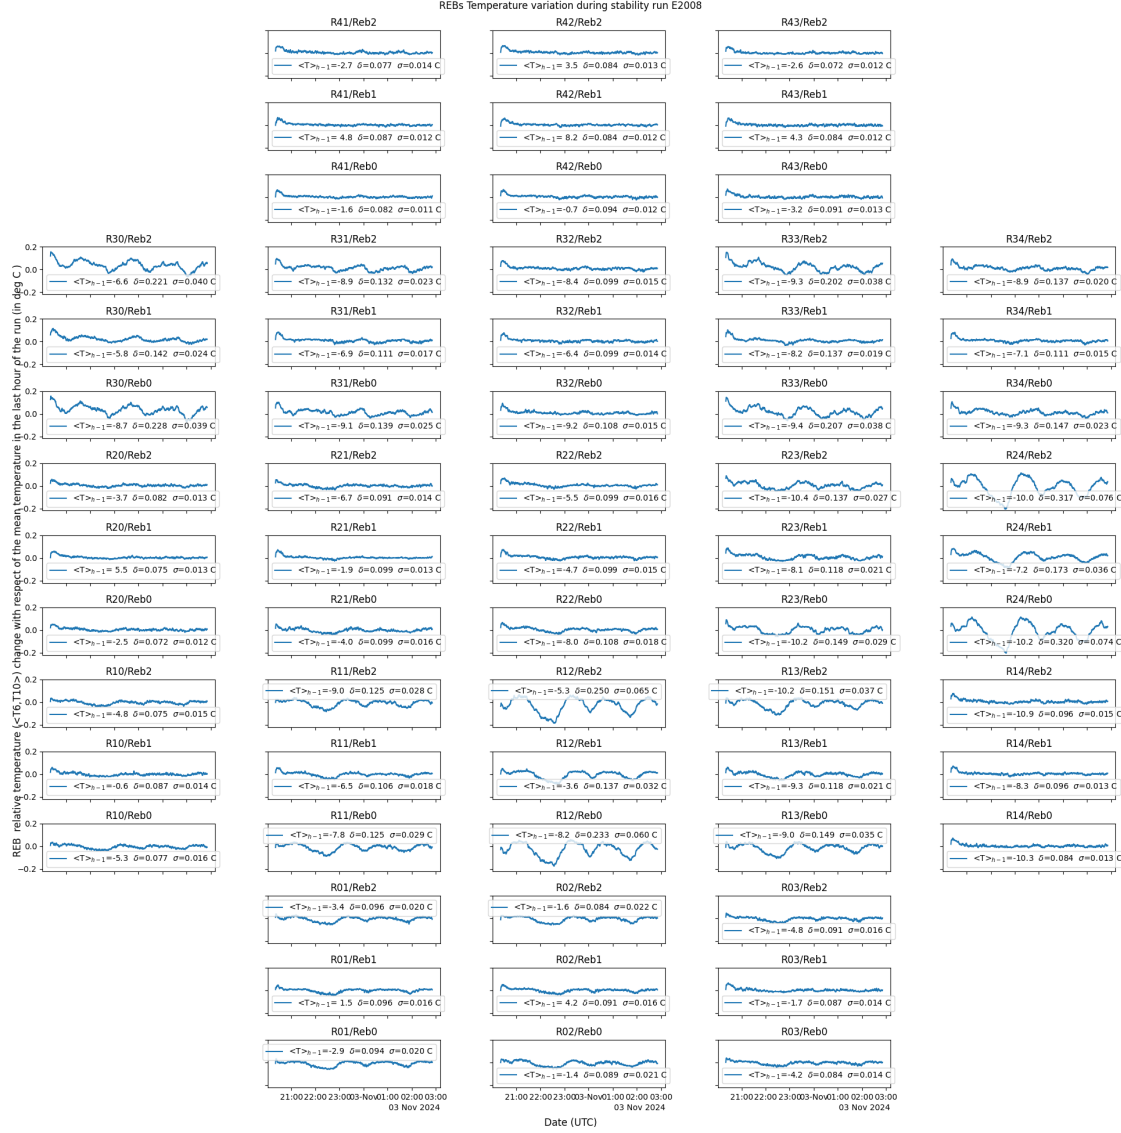


Figure 114: REBs temperature change during run E2008 (IDLE_FLUSH Off). In the plot the temperature of each REB is relative to the mean temperature of the last hour of the run. For each Reb the value of this mean temperature ($\langle T \rangle_{h-1}$), the peak-to-peak temperature change (δ) and the temperature dispersion over the run (σ) is given. We see clearly that not all temperature variations are the same, for example the REBs of R24 have larger variation than those of R14

All those temperature changes and non-uniformity are small: at maximum we observe a temperature change of 0.6 deg in Run E1496, which will induce a gain change of $\sim 4 \times 10^{-4}$. Also running the focal plane with more stable temperature over time, e.g., by not using IDLE_FLUSH (in run E2008, temperature changes induces a gain change of at most $\sim 2 \times 10^{-4}$), will further




Figure 115: Ratio of the measured signal to the mean of all channels of R24 relative to that for R14: (left) without temperature correction; (right) with temperature correction. This confirms that there is a dispersion in temperature that induces a dispersion in gain response. The crude gain correction (6.5×10^{-4} per deg C) based on the measured REB temperature, essentially fixes the initial dispersion observed.

improve the gain stability. Note that on this front we have a plan to further uniformize over time the REB usage (mimicking exposure readout, even when not taking an exposure). Still it is interesting to notice that the temperature change over the focal plane has a dispersion for the REB electronics, with larger temperature variation always in the same REBs (e.g., R24/Reb2). At this stage the origin of this non-uniformity is unknown. In Figure 115, we underline this non-uniformity by comparing the average signal measurement in R24 (large temperature variation observed) and R14 (minimal temperature variation) during run E1496 (stable temperature). A crude temperature correction (we did apply a gain correction of 6.5×10^{-4} per deg C to all amplifiers based on their REB temperature) brings the agreement of response between the two rafts to better than 1×10^{-4} .

6.9.3 ITL gain glitches

Following studies done in Run 6, we know that subtle clock phase changes between REB, associated with the current DAQ clock scheme, induce from exposure to exposure changes in bias levels and gain, in ITL sensors only. We observed from exposure to exposure, bias shifts of a few ADU and gain change in the worst case of 0.08% (peak to peak). For ITL sensors the system response aggregates in what we call “families”: a given amplifier having its gain changing among a limited set of possible values. We do observe the same effect in Run 7 (see Fig. 116). Notice that the channels with the largest gain changes observed in Run 7 (in run E1496 for this study), are not the same as in the runs we studied in Run 6. As the timing de-synchronization between the REBs, will change each time the DAQ is re-started, we do expect such differences.

Based on the studies done for Run 6, this effect will be extremely difficult to correct in the data. Even if it is small ($< 10^{-3}$), we do explore a solution at the level of the DAQ to avoid it. The main interest of such a hardware fix could be in the resulting stability of the 2D BIAS shape, more than in avoiding the gain glitches, which are small.




Figure 116: The measured signal for R01_S12_C01 is shown for the various exposures of Run E1496. The signal is normalized to its average value for the run, after applying a temperature correction and removing common-mode signal variations using the mean signal from the e2v sensors. This ITL channel exhibits the largest gain glitch in this run, with a peak-to-peak variation of 0.08%. We observe three distinct gain levels (“families”) for this channel, which we attribute to different DAQ synchronization states between the REBs of this raft.

6.10 Guider operation

Beginning with Run 7, it became common to run the corner rafts in their nominal (guiding) configuration while acquiring the EO data for the rest of the focal plane. This section describes guider-specific acquisitions designed to verify guider requirements and measure performance. For a complete description of guider requirement verification, please refer to LCA-20583.

6.10.1 Guide mode

Because different regions of pixels on each GREB sensor can be read out, the guider requires a version of GREB firmware that implements a separate sequencer for each sensor. When in science or full frame mode, there is only one sequencer in the GREB. However, because the GREB can be configured to read out a different portion of each sensor, special firmware must be loaded that contains a sequencer for each sensor. Additionally, the DAQ embedded processor connected to the GREB must also run special guide software and the DAQ synchronous timing system must be configured to allow the separation of readout commands between the guider and the rest of the focal plane. Thus, switching between guide and science modes is a non-trivial operation.

Switching from science to guide mode has several steps:

1. Power off HV bias and sensor power
2. Reboot the GREB into the guider (multi-sequencer) firmware
3. Reconfigure the DAQ and CCS
4. Power on the sensors and HV bias

Currently, the GREBs will boot into the science firmware on power up, but we are expecting to change this default to guider firmware prior to operations. Similarly, though powering the HV bias off will always be necessary when rebooting the GREB, as it is controlled by the GREB firmware, but powering off the sensor was done out of an abundance of caution. If it becomes common to switch back and forth, which is not expected, it may be worth investigating whether this is actually necessary.

6.10.2 Guider timing

The guider is designed to operate in a continuous loop, alternating between integration and readout. A set of integration and readout makes an image, a stamp. To first order, the time between stamps is the sum of the integration time and the sequencer execution time. The sequencer execution time varies with the ROI size in a way that is fixed for a version of the sequencer program, since the ROI size determines the number of parallel/serial transfers. However, due to the details of the DAQ implementation, there is also a contribution from transporting the data within the embedded DAQ processors. Figure 117 shows the inter-stamp timing for all guide-specific runs in the nominal ROI configuration. The beginning of the readout of each sensor is synchronized to \pm one system clock cycle (10 ns) using the same mechanism as for science readout.




Figure 117: Inter-stamp time for nominal ROI

Noise Study Configurations For any guider configuration, the ROI size is the same for all guide sensors, but its location on each sensor varies. This means that the total readout time is the same, but at any given time within that total, any given sensor can be fast-flushing rows, flushing columns, or reading out pixels. To examine the sensitivity of the GREB to noise

induced by the phasing of readout and clearing, a series of configurations was defined. Beginning with a single sensor as the baseline, two sensors on the same GREB were configured to read out with their ROI locations the same (*aligned*) and not overlapping (*unaligned*.) This was to measure whether noise can be induced on sensors within the same GREB. This was followed by four aligned ROIs, one sensor on each GREB and four sensors all unaligned, to check for noise between GREBs. Finally, with an ROI defined on each of the eight sensors all aligned and all unaligned between GREBs.

ROI Study Configurations We also collected a set of data with ROIs of different sizes and with different integration times to measure the guiding rate and noise. These were single sensor configurations to be compared with the nominal size baseline from the noise study configurations. ROIs can also be specified to span sensor segment boundaries, and so a configuration was defined for that.

6.10.3 Noise investigation

We measure the noise level of ROIs acquired under various configurations, shown in Table 11. We took 20 images in each configuration, where each image is acquired over a 15 s exposure time. Due to different ROI sizes (and thus different read-out frequencies), the number of frames within each image varies. The noise is calculated as the standard deviation of the entire ROI from R00_SG0, and averaged over all frames from all of the 20 images. For the split ROI (last row in Table 11), the noise is measured on the left and right half of the ROI respectively. The images were taken on 30 Nov. 2024 and 01 Dec. 2024. We note that all images taken on 30 Nov. 2024 suffer from an abnormally high-gain sensor state, where the counts level in each image is about one-tenth of the expected values (Section 8.8). This affects most of the rows in Table 11 except the last two rows. The cause of such an abnormal state is under active investigation. Regardless of the issue, an increase in noise is seen when ROIs are unaligned on a single GREB, but not among GREBs.

Table 11: Summary of results for the different Guider configurations.

ROI Size	Integration Time (ms)	Number of Sensors	Number of Rafts	ROI Alignment	Rate (Hz)	Noise (ADU)
Noise Study Configurations						

Table 11: Summary of results for the different Guider configurations.

ROI Size	Integration Time (ms)	Number of Sensors	Number of Rafts	ROI Alignment	Rate (Hz)	Noise (ADU)
50×50	50	1	1	n/a	9.28	5.60
50×50	50	2	1	aligned	9.27	5.64
50×50	50	2	1	unaligned	9.26	8.63
50×50	50	4	4	aligned	9.26	5.61
50×50	50	4	4	unaligned	9.26	5.64
50×50	50	8	4	aligned	9.23	5.65
50×50	50	8	4	unaligned	9.23	5.68
Nominal Configurations						
50×50	50	8	4	aligned	9.22	5.65
50×50	50	8	4	unaligned	9.23	8.67
ROI Study Configurations						
400×400	200	1	1	n/a	1.67	4.03
400×400	50	1	1	n/a	2.23	3.95
400×400	5	1	1	n/a	2.48	3.91
10×10	50	1	1	n/a	11.80	13.56
400×400	50	1	1	SplitROI	2.23	7.24, 7.09

6.10.4 Impact on science sensors

We compare two runs, E1110 (guide sensors in imaging mode) and E1290 (guide sensors in guider mode), to study the impact on science sensors from running guide sensors in guider mode. Figure 118 demonstrates that the read noise is mostly consistent between the two runs. The only exception is R01_S02 on which a few channels have slightly lower read noise in E1290, and R01_S11 C00 which has a higher noise in E1290. Figure 119 shows the read noise comparison of R01 as well as the temperature measurement on S02. E1290 has a slightly lower temperature, which might explain the lower read noise. Notice that this deviation is also seen in the read noise comparison with the initial run (Section 4.4.3) as well as the PTC noise (Section 4.4.4), which indicates that it is not caused by the guiding mode.

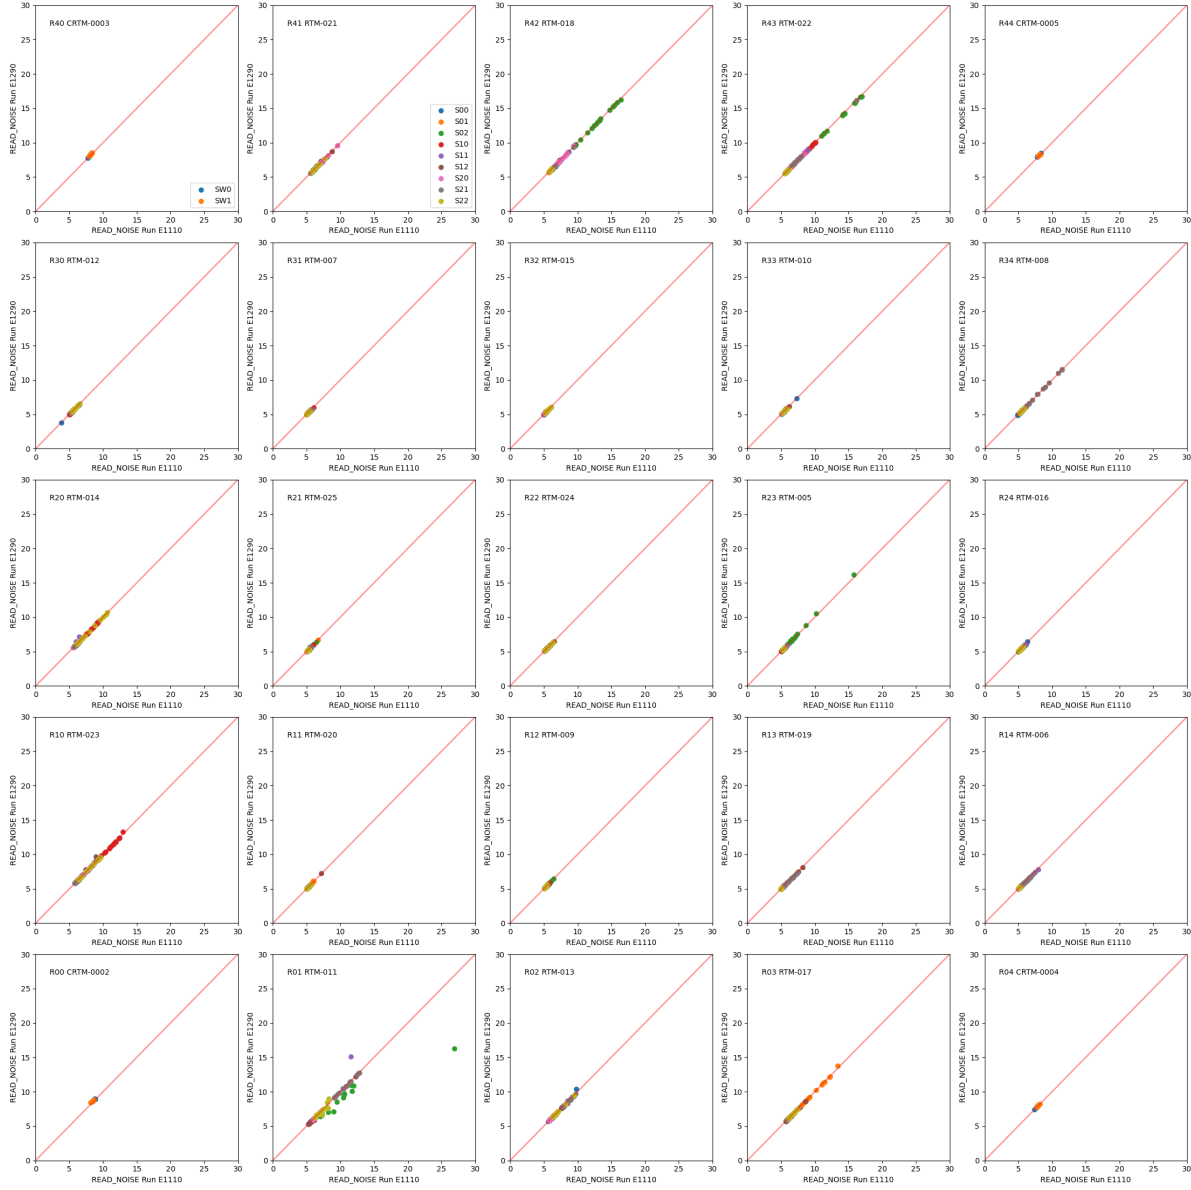


Figure 118: Impact on science sensors read noise from guide sensors running in guider mode. E1110 has the guide sensors in imaging mode, and E1290 has the guide sensors in guiding mode. The two runs have consistent read noise except for R01_S02, where E1290 has slightly lower noises. This might be caused by a temperature fluctuation (see Fig 119).




Figure 119: Left: Read noise of R01 compared between run E1110 (guiders in imaging mode) and E1290 (guiders in guiding mode). Sensor S02 has lower read noises in E1290, and S11_C00 has a higher noise in E1290. Right: Temperature measurements from R01_S02 in E1110 and E1290. Run E1290 has a slightly lower temperature.

6.11 Summary

In summary, our comprehensive characterization of the performance and stability of LSSTCam reveals several key findings. The comparison of initial and final Run 7 camera metrics on Cerro Pachón, using standard B protocol and dense red PTC data sets, highlights the consistency and stability of serial and parallel CTI measurements. Dark current measurements remain stable across the focal plane, with notable improvements in certain rafts due to effective light leak mitigation. Bright defects show close agreement between runs, with no significant new defects.

Evaluations of linearity and PTC turnoff from PTC runs indicate a decrease in full-well capacity for e2v sensors, while PTC gain shows a minor increase. Read noise remains consistent with Run 6 for both e2v and ITL sensors, and PTC noise does not exhibit significant deviations. The brighter-fatter coefficient increases for e2v sensors due to lower parallel swing, and the brighter-fatter correlation remains consistent. Row-means variance shows close agreement, with a slight decrease for e2v sensors. Divisadero tearing experiences a significant decrease for e2v sensors in the final condition. Dark defects remain consistent between runs, affecting a very small fraction of pixels.

A significant achievement is the substantial decrease in persistence signal for e2v sensors, attributed to the lower parallel swing. This optimization marks a notable improvement in the camera's performance.

Differences between initial and final measurements emphasize minimized persistence in e2v sensors, decreased full-well capacity, increased brighter-fatter strength, and reduced divisadero tearing.

Problematic amplifiers are listed, classified as high noise or dead/limited signal, with insights into changes in amplifier status during testing periods.

Defect stability is another critical aspect, with bright defects found to be stable and rare, while dark defects are more dynamic and abundant, showing significant temporal dependence.

Bias stability is explored, revealing large variations in ITL bias jumps and residual 2D shapes in e2v yellow corners after correction. These instabilities depend on acquisition sequences and exposure times. Further studies are required to converge on the best mitigation strategy for

the start of the LSST survey.

The glow search identifies glowing regions in long dark images, noting some persistent issues. The illumination corrected flat section assesses gain matching and relative QE using a smooth model fit to illumination patterns.

Full-well measurements are evaluated to set SAT and SUSPECT pixel masks, comparing spot images to flat illumination data. Nonlinearity studies model the nonlinearity using spline functions, with results deemed sufficient for the camera's needs.

Gain stability is scrutinized, with temperature dependencies and potential issues with local amplifier changes noted. Maintaining stable conditions is crucial to ensure consistent gain.

We describe guider sensor operation, focusing on guider-specific acquisitions and performance verification. Noise investigations reveal the impact of various configurations on noise levels, and we provide a comparison to science sensors when guide sensors are operated in imaging mode.

7 Sensor Features

7.1 Tree rings

Tree rings are concentric variations in silicon doping concentration the effects of which can be observed in flat images. The tree rings are characterized in Park et al. (2017, 2020) and the impact of the tree rings is assessed in Esteves et al. (2023). In this section we describe measurements of the tree rings for each LSSTCam sensor from the test data taken in Run 6 and Run 7.

7.1.1 Centers of the tree rings

From the characterization study cited above, the center of the tree ring pattern is known to have four possible positions with respect to the center of a sensor. This is because four CCDs are cut from one wafer. To date we have been using the four different average positions for each center of the tree ring pattern, according to the pattern direction, because measuring the tree ring pattern for each sensor is difficult due to their low amplitude, and low contrast with respect to other features such as the “CMB” pattern. However, in Run 7 we obtained new data with 0V of back bias voltage with the diffuser installed (Sec. 4.6 and Fig. 3), which increases the amplitude of the tree ring pattern in the images, allowing us to revisit the measurement of each individual center (E1050 for red, E1052 for blue, E1055 for nm850, E1056 for nm960 LEDs).

Figure 120 shows the positions of the tree ring centers measured for the 189 science sensors. The measurements are concentrated around four averaged positions. However, as now we have better individual measurements, we decided to use the centers as evaluated individually for each sensor instead of the average value based on which quadrant of the wafer a sensor came from.

7.1.2 Radial study

We performed a radial study of the tree ring pattern to see if the rings are perfectly circular in shape.

Figure 121 illustrates the transformation of a flat image into a radial profile plot with the x axis




Figure 120: The centers of the tree ring patterns measured for all 189 LSSTCam science sensors. Red points indicate the average center for each direction.

being the distance from the center of the rings.

7.1.3 Effect of diffuser

We expected that with the diffuser installed, contributions from effects such as CMB and weather patterns discussed in Section 3.1 would be diminished. Comparing results for R22_S12 from Run 6 run 13379 (without diffuser) with Run 7 E937 (with diffuser), we verified the significant improvement from use of the diffuser (cf. Fig. 123 and Fig. 124).




Figure 121: A flipped image on diagonal line from the center of the ring, and subtracting from each other.




Figure 122: Radial study of the Tree Rings. Right: image subtracting left to right, right to left.

7.1.4 Voltage dependency

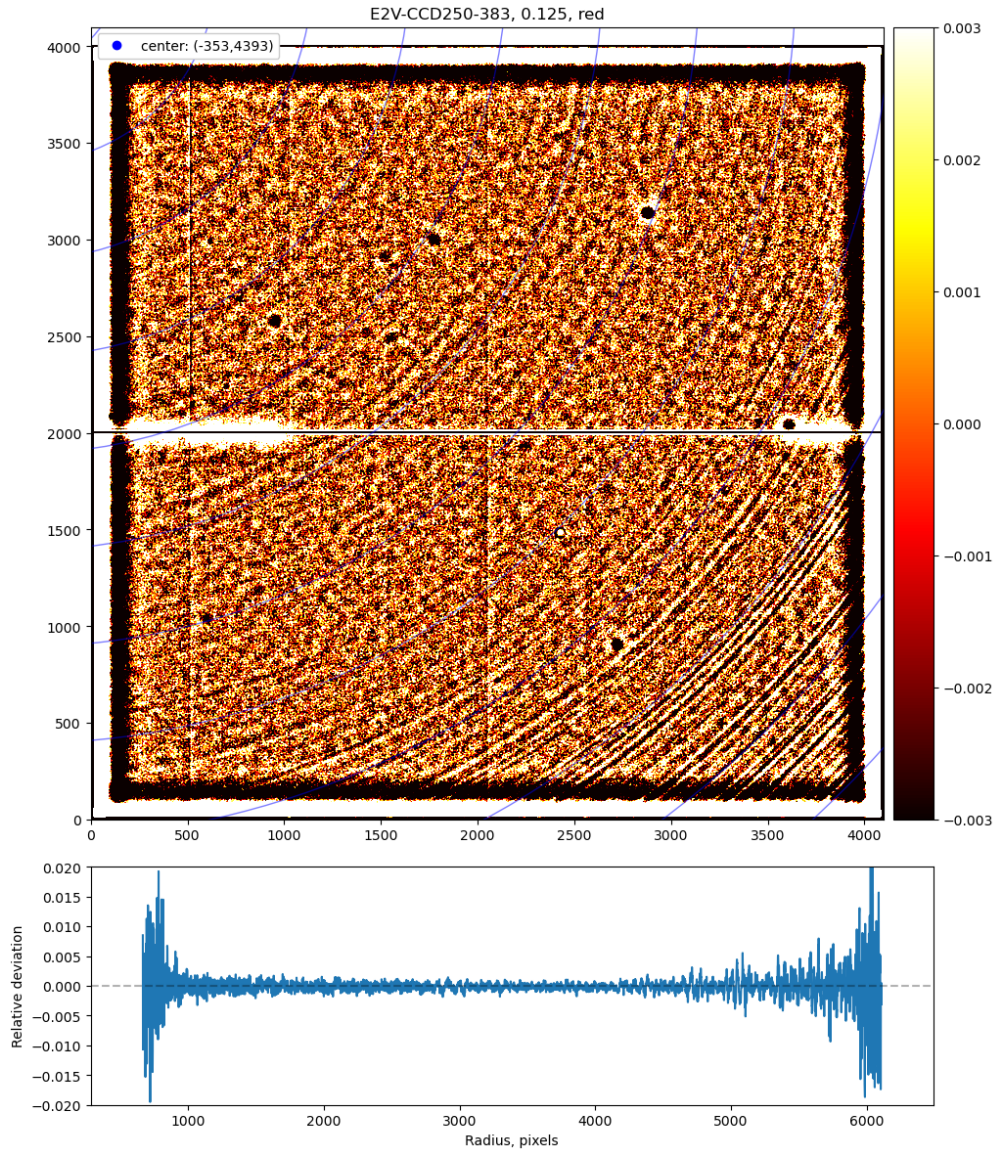


Figure 123: Tree ring without diffuser

7.1.5 Wavelength dependence

The tree ring effect does not depend much on wavelength, as studied in (Park et al., 2017, 2020). However, when the flat images for different wavelengths are compared, we can see the other sensor effects dominate in shorter wavelength (laser annealing) and longer wavelength (fringing) over the tree ring patterns. Figure 126 shows the tree rings in red and blue flat images, and Figure 127 shows the difference image between red and blue to highlight the differences. We can see that the laser annealing effect is left since the amplitudes of the tree

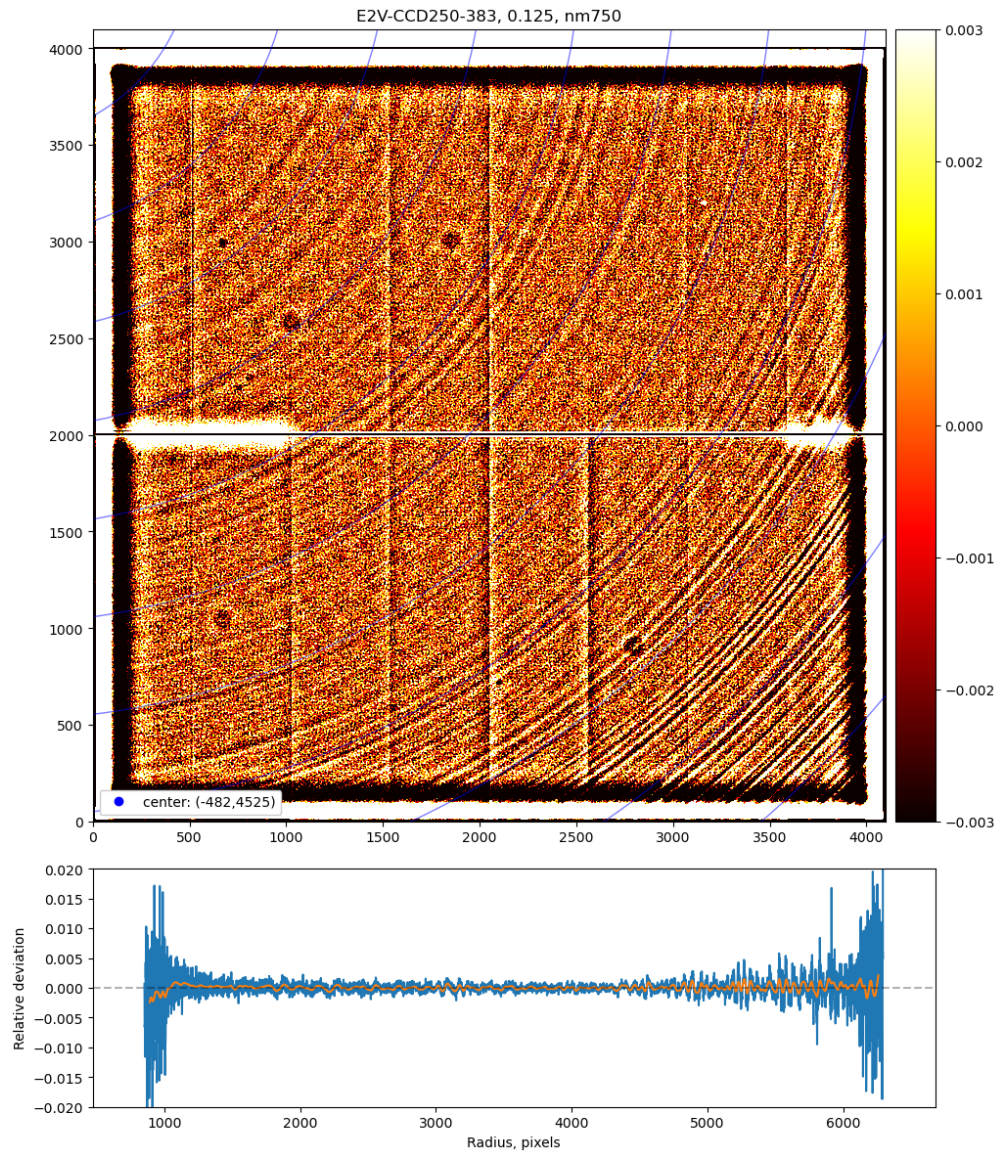


Figure 124: Tree ring with diffuser

ring effect in blue and red images are similar.

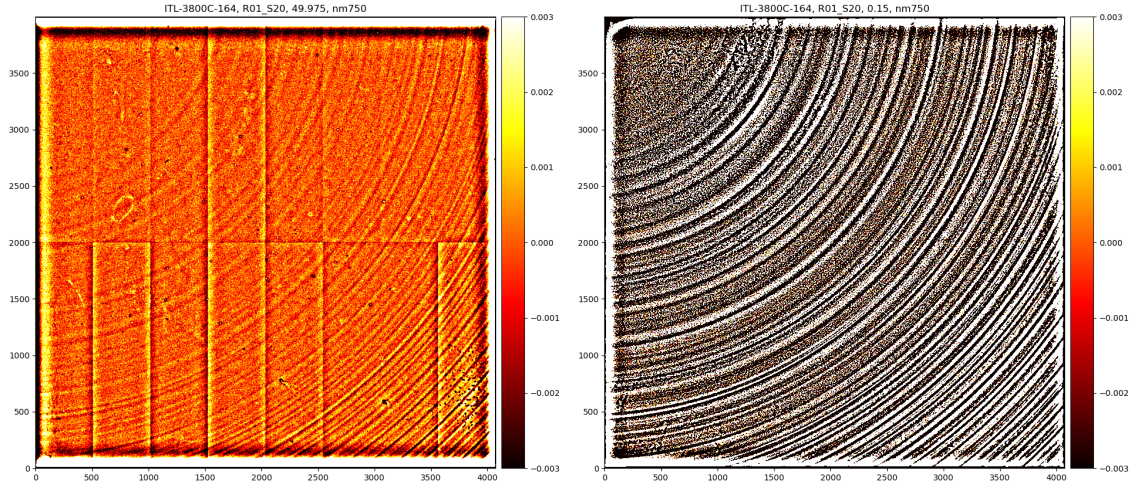


Figure 125: Comparison of the tree ring patterns on the sensor R01_S20 with (left) and without (right) back bias voltage (50 V) applied. The applied back bias voltage clearly reduces the impact of the tree ring effect.

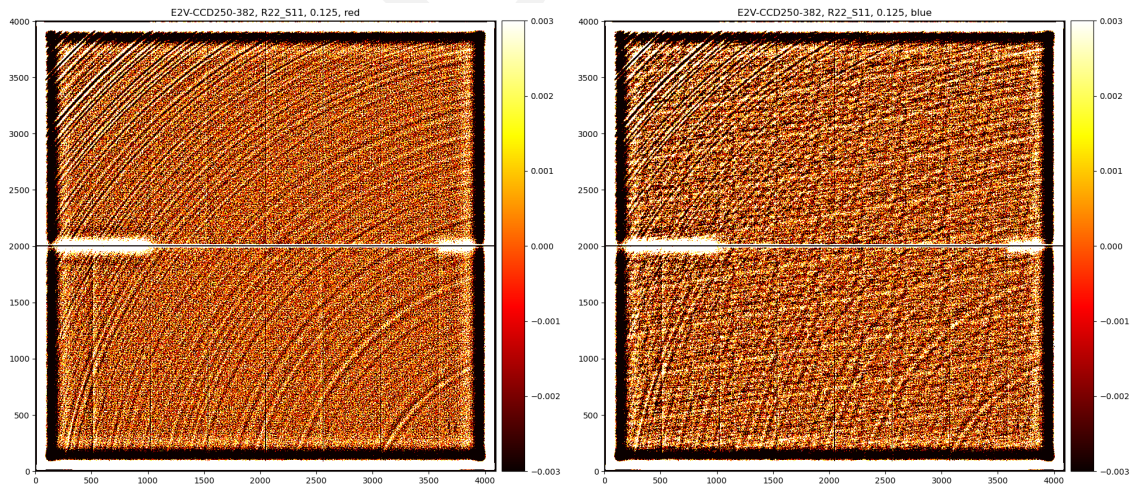


Figure 126: Comparison of the tree ring patterns on the sensor R01_S20 for red (run E1050, left image) and blue (run E1052, right image) wavelengths, without back bias voltage applied.




Figure 127: Difference between the blue image and the red image, highlighting the laser annealing pattern that is evident in the blue image but not in the red.

7.2 ITL dips

One of the phenomena that was studied in the later part of Run 7 was so-called ‘ITL dips’. These were discovered in LSST ComCam on-sky data as bleed trails from bright stars that traversed the entire detector, crossing the amplifier boundaries. These bleed trails are unique though in that the core of the bleed trail is actually ‘dark’ compared to the wings of the trail, with a flux $\sim 2\%$ less relative to the rest of the bleed trail, as seen in the profile in Figure 128.

A possible explanation for these ‘dark dips’ involves the depletion of holes from within the channel stop implants between columns, similarly to the divisadero effect at the segment edges of e2v sensors. The lower hole density between those columns creates a local lateral distortion of the drift field that shifts a fraction of the signal away from the center columns and towards the outer columns. When a highly saturated star is bleeding, electrons are flowing across potential barriers in both the column and row directions. It is possible that, with a sufficient influx, some electrons recombine with the holes in the channel stops. In ITL sensors, the channel stop implants are continuous across the whole sensor, so it is possible for this lower hole density to propagate away from the saturated star, both above and below the source of electrons.




Figure 128: Profiles of an ‘ITL dip’ in an on-sky frame of sensor S22 of ComCam. The profiles are extracted by computing the medians over the columns for the parts of the frame both above and below the saturated star. The blue vertical lines mark the extent of the saturated footprint across columns.

We investigated whether ITL dips could also be observed in the CCDs of LSSTCam. For this study, we used spots and rectangles projected onto the focal plane by the 4K projector. The

spots were approximately 30 pixels across and were projected onto every amplifier segment of each detector. The rectangles were only in the top right amplifier (C10). One consideration with this spot projection was that the projector also provided background illumination. This led to the spots having a peak signal only 6 times greater than the background and the rectangles having a peak signal 30 times greater than the background.

We were unable to find any evidence of ITL dips in the either circular or rectangle spot images. For the rectangular spots there were two tests done, they both looked for any dips in the neighboring amplifier with different binning strategies. Example results of the two rectangular studies are shown in Figure 129 using just a base average of the rows and Figure 130 shows the results of binning using background corrected and normalized binning schema. For the circular spots, the slices were done much closer to the spot, only 200 pixels away from the saturated spot. Figure 131 shows example cutouts from the circular spot images. In all three cases, these cutouts show the background pattern of the projector, but no 2% dip.

While we were not able to find evidence of the ITL dip in Run 7 data, it is still not clear whether the effect will be visible in LSSTCam on-sky data. Study of the effect on ComCam sensors showed that not all CCDs were affected to the same extent, ranging from one sensor that was particularly sensitive, to two that were almost not affected. It is possible that the contrast and brightness that we could achieve in the lab was not representative of the on-sky measurements. But the ComCam sensors were run with a different sequencer file, and potentially the different sequencer files could explain the effect. Also, the CCD operating temperature of ComCam, which was $\sim -90^{\circ}\text{C}$ or higher, was higher than -100°C for LSSTCam. It is not clear if this is related or not.

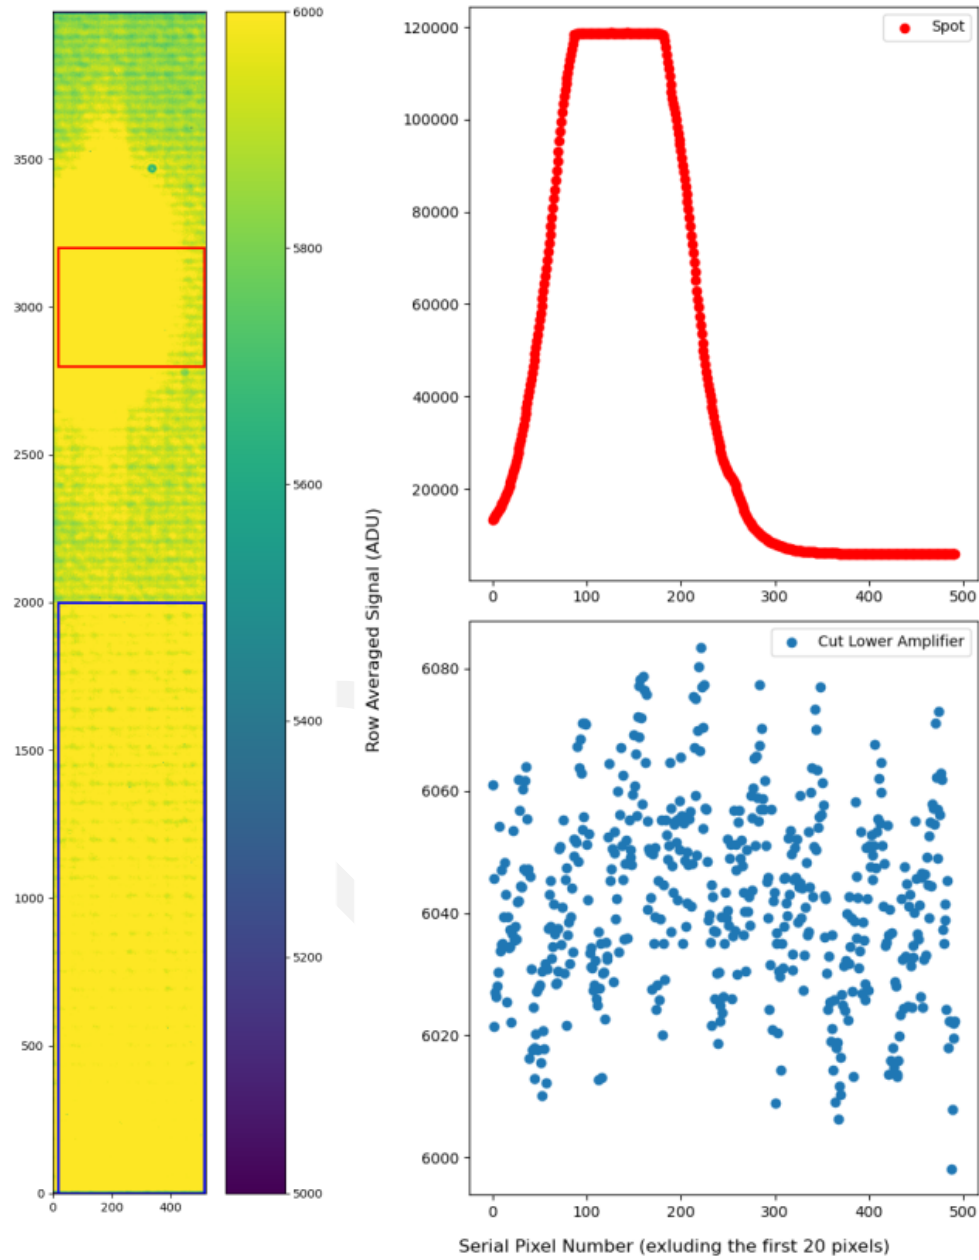


Figure 129: (Left) Cutout showing two of the amplifiers of the one of the images with a saturated rectangle 'spot' from the 4k projector. The red box shows the area that was used to highlight the saturated area while the blue box shows the area below that was used to look for the ITL dip. (Right) The row averaged pixel counts in ADU for the saturated region and the region below. If there was an ITL dip, we would expect to see a 2-5% dip around the saturated region in the amplifier below it.




Figure 130: An example of a slope-corrected, normalized mean of the columns in C00, the amplifier below the rectangular spot, of both a heavily saturated spot image and an unsaturated spot. The black lines signify where we would expect the ITL Dip to occur if present.

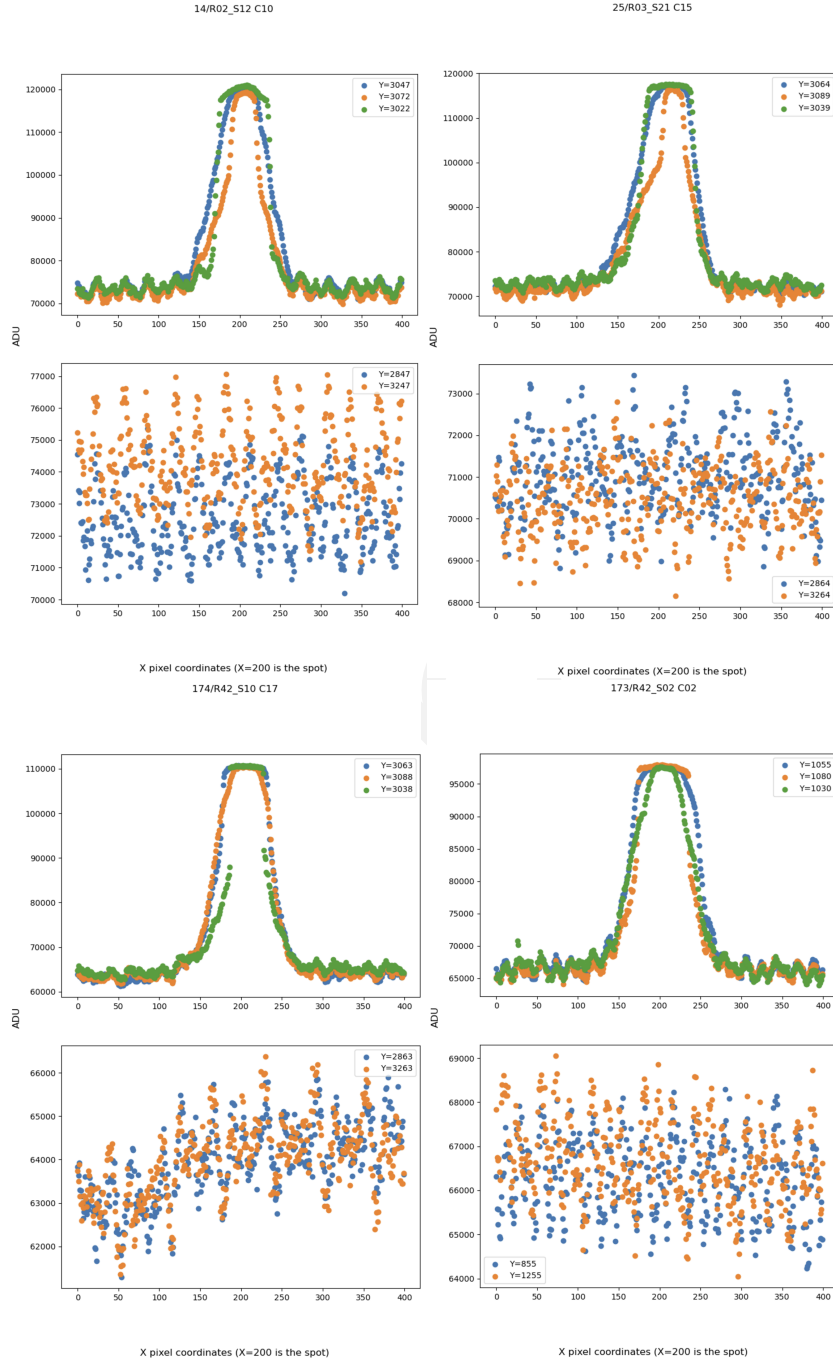


Figure 131: Different cutouts using circular spot data. The titles show the detector number, detector name, and amplifier name. The top plots are horizontal cutouts across amplifiers centered and around the spot. The bottom plots are horizontal cutouts about 200 pixels away from the center of the spot. None of the bottom plots show a 2% dip around the area of the spot.

7.3 Vampire pixels

A category of sensor feature found on some ITL sensors, that has recently benefited from fresh attention, now has a new name. They are called *vampire pixels* because of their curious flat field response: a group of pixels with photo-response exceeding the flat-field mean, surrounded by a concentric distribution of pixels with photo-response below the same flat-field mean. The *vampire pixel* name sticks because the over-responsive pixels have apparently “sucked” signal from the rightful receivers, a sort of *reverse brighter-fatter* effect excited simply with flat-field illumination.

The sizes of these *vampire pixel* complexes can typically extend to tens of pixels in radius, which place strong constraints on their origin. Also, it turns out that all prominent *vampire pixel* complexes are also seen in their (ITL) phosphorescence response which is indicative of the backside surface layer (cf. Sec. 7.4). This means that *vampire pixels* are likely to appear also in dark images, but only if trigger illumination is delivered a few tens of seconds prior to the dark acquisition (cf. vampire transients, Tab. 14).

All known cases appear to round or with circular symmetry. There are plenty of cases with similar pixel complexes that lack the central group of pixels with *high-amplitude* photo-response excess, but that the photo-response deviation is simply divided into a larger number of pixels or in locations further from the center of the complex. We suggest that the underlying origin of these is common with the easier-to-detect *bright pixels* (cf. Sec. 7.3.1) but appear with different response properties simply because of mundane geometric details. Different detection algorithms may therefore be required finding those *vampire pixel* complexes that do *not* show central bright pixels as opposed those that *do*. Moving forward, we choose not to invent a new name for the former type, but refer to them all as *vampire pixel* complexes.

7.3.1 First observations

Initial identification of these on ComCam may have been in a study that called them *bright pixels* by A. Roodman (20240827) and quantified in more depth by A. Ferté in a ComCam defects study (20241112). First electrostatic simulations performed to help understand them were made by C. Lage (20241119) who inferred that a circumferential surface charge variation⁵ on the backside electrode could reproduce the sort of charge redistribution observed

⁵We suggest that any such variations would necessarily require that the backside electrode has ceased to act as an electrostatic equipotential locally, as it does elsewhere on the sensor (with total surface charge density

(while conserving photo-conversion charge) – and so these may be effectively described as pixel boundary distortions throughout, mediated simply by lateral (non-axial) contributions to the drift field. Any such lateral fields would mean a localized loss in pixel fidelity, not limited to the sensor’s thickness scale (10 pix) as are apparently relevant in other known pixel distortion mechanisms (brighter-fatter effect, tree rings, edge rolloff, tearing, pixel boundary distortions due to midline implant, and hot columns).

Since ComCam on-sky data became available, more attention has been paid to these features and how they may impact source detection and photometric determination of nearby field sources. As a direct consequence of tighter device screening and selection, *vampire pixels* are less common on average in the 88 ITL sensors of LSSTCam than they are in the nine sensors of ComCam.

7.3.2 LSST Camera vampire pixel features

One prominent example of such a feature is located in the LSSTCam focal plane on R01_S00_C13-4. This feature is often overshadowed by the bright, dark current “scratch” in close proximity (only when HV Bias is on). In Figure 132 and Table 12 we include this example along with two other prominent *vampire pixel* complexes (and others) located on ITL sensors in LSSTCam to describe their individual properties. Inspection reveals a broad parameter space describing these pixel complexes that can cause distortion in one way or another as soon as they are used to record cosmic sources: astrometric and shape transfer errors inferred from flat response, and background estimation or source confusion errors from the pixel complexes’ phosphorescence properties.

The pixel complexes evaluated in Table 12 were chosen based on their proximity to the prominent *vampire pixel* appearing at the center of the corresponding image given in Figure 132, and include (on average) 3 other, nearby complexes that may be more representative of these artifacts found on ITL sensors in the Main Camera focal-plane. In the Table they are indicated by their name (“babyX”) followed by the clocking angle where they can be identified relative to the prominent pixel complex located at the center. From this list, it appears that the *vampire pixel* complexes may be reliably identified by applying a OR combination of thresholds: under-response less than 90%, **or** an over-response exceeding 120% and some consideration of the presence of phosphorescence. The phosphorescence, more than anything, may help

governed only by the normal component of the electric field strength within the silicon, responding to the HV Bias potential, and so on).




Figure 132: Three prototypical *vampire pixel* complexes occurring in LSSTCam. Each of these have counterparts that appear in phosphorescence (transient dark). These are each described in Table 12. In R01_S00, Two “baby” vampires appear in the 8:00 and 8:30 positions that each possess 300–400% nominal flat response. The recorded phosphorescent counterparts for R03_S10 and R20_S20 are dependent on HV Bias (cf. Figs. 137a through 137d) and should provide some constraints as to the pixel astrometric shifts at play near these *vampire pixel* complexes.

Table 12: Sample vampire pixel complex parameters. The last column “P” indicates whether a concentric phosphorescence center is present.

<i>vampire pixel</i> complex	radius at 99% flat response	minimum under-response	maximum over-response	peak at center?	P?
R01_S00_C13-4	200 pix	1.4%	1660%	distributed (10 pix offset)	yes
baby1 (8:00)	6 pix	83%	375%	yes	yes
baby2 (8:30)	4 pix	84%	290%	yes	yes
R03_S10_C15	36 pix	21%	1570%	yes	yes
baby3 (8:30)	4 pix?	97%	152%	yes	no
baby4 (10:30)	4 pix?	98%	120%	yes	no
baby5 (4:00)	4 pix?	89%	NA	no peak	yes
R20_S20_C13	52 pix	40%	829%	no (ring-like)	yes
baby6 (10:00)	NA	NA	108%	yes	no
baby7 (3:00)	NA	NA	119%	yes	no
baby8 (7:00)	NA	NA	207%	yes	no

to distinguish the dark pixel complexes *without* central bright pixels from dust spots (which presumably would not preserve flux).

A handful of dust spots are seen in these images that were not included in Table 12. They would presumably be detected as *dark pixels*, provided the lower thresholds are raised to levels appropriate for their detection.

A listing of ITL sensors in the Main Camera focal-plane showing such *vampire pixel* complexes is given in a table in Section 7.4, Table 14). Some fraction of the spot-like transient features counted on each sensor are associated with coincident reductions or enhancements in the flat-field response, consistent with the above descriptions of *vampire pixels*. The list shows that 83 of the 88 Main Camera ITL sensors contain finite numbers of potential *vampire pixel* complexes, typically fewer than 30, revealed by spot-like phosphorescence. Significant coincidence is expected with pixel complexes that show flat-field response variations (some of which would be simply due to sensor surface hosted dust spots (*dark pixels* due to absorption and/or shadows, but sometimes appear as *bright pixels* when redirection due to reflection occurs). We therefore focus on coincidences between features showing flat field response variations and transient, phosphorescence features and identify them as having a non-dust origin. Certainly, e2v sensor surfaces also show dust-related features, but only ITL sensors appear to exhibit *phosphorescence* as we describe here.

Based on this listing, a total of 17 ITL sensors in the Main Camera host more than 30 *vampire pixel* complexes each. These identifications may be used to study in depth more fully the science impact of their transverse electric fields as well as their phosphorescent properties (when recording is preceded directly by illumination by a star in the previous image).

As a proof of concept, a task was added to `eo_pipe` to search for bright defect pixels in combined flats. Figure 133 displays the resulting distribution, which efficiently identifies the ITL set of sensors. Without looking more closely at the specific regions flagged, distinguishing the *vampire pixel* complexes from *reflecting particulates* would be impossible. Conversely, repeating this task to identify dark pixels with a threshold of 0.90 (90% of flat level), we expect to see a combination of (flux conserving) *vampire pixel* complexes and garden variety (flux attenuating) dust to appear. (Currently, we do not know whether dust particulates preferentially stick to e2v or to ITL sensors.)




Figure 133: Results of the `eo_pipe` task written to search for bright defects in combined flats. Using a threshold of 1.2 (120% of flat level), the findings highlight the 8 RTMs that operate ITL sensors (as well as those CRTMs with sensors operating in science mode).

7.4 Phosphorescence

The Run 7 *persistence* optimization process (cf. Sec. 5.1.1) used a short EO image acquisition sequence and analysis script, which rapidly provided persistence performance metrics as feedback for each configuration tested. Thus, as soon as the e2v sensors were shown to be nearly free of *their* undesirable effects by reducing their clock swing voltages from 9.3V down to 8.0V (Sec. 5.1), a similar persistence (or memory effect) was immediately noticed, affecting a subset of the ITL sensors. This late discovery gained immediate interest for at least two reasons: (1) that it had not been detected in prior EO campaigns, and (2) that the new memory effect on certain ITL sensors was morphologically distinct from what had just been *cured* on the e2v sensors.

The ITL sensors with the largest memory effect were evaluated, and the following observations were made:

1. The morphology of the expressed memory effect in the first dark image acquired after the trigger (the saturation flat) was reminiscent of the “*coffee stains*” seen on the same sensors in flat field response, but with the opposite polarity. The “*coffee stains*” are commonly assumed to be associated with minor, localized variations in the sensors’ antireflective coatings or perhaps a very thin, dead layer associated with the backside surface: they tend to be larger in amplitude when shorter wavelengths are used to expose the sensors with flat field illumination.
2. The attenuation timescale of the memory effect is curiously comparable to the timescales that were seen in the persistence suffered by the e2v sensors (which are believed due to exposure of surface states by the collected conversions, on the semiconductor-insulator interface on the front side): exponential time constants of between 20 and 40 s, which curiously are in turn very close to the nominal exposure cadence for the LSST survey.
3. The similarity in memory effect time constants (de-trapping charges from surface states near the channel on the front side – the e2v case – vs. either de-trapping of charges located near the backside window surface or relaxation by photon emission by some excited states there – the ITL case) can be thought to favor the electron de-trapping mechanism, if only from the other surface. Otherwise, the nearly matched time constants would have to be seen as an improbable coincidence.
4. A list of 12 ITL sensor serial numbers corresponding to those showing the memory ef-

fect was communicated to Mike Lesser at ITL. The list of parts shared certain properties according to his notes, and led him to develop a placeholder theory that would partially explain the mechanism. If true, it could explain what might be responsible for both the coffee stains and the memory effect with similar spatial distribution. He wrote that he tried, but was unsuccessful in diagnosing, using optical characterization tools (e.g., ellipsometer), any changes in optical constants on the affected regions of the “stained” sensors. The origin of the “stains”, according to this theory, is as a consequence of there being “raised spots” on the sensors’ backside surfaces that survive the final silicon acid etch. The raised silicon areas could potentially be trapping the resist used during the cleaning process that directly follows the etching step. Lesser wrote that the resist is wax-based and *does* fluoresce. If the theory is correct, he suggests that the medium would definitely be located *under* the AR coating and related neither to the coating nor the oxidation processes.

5. Discussions among the Rubin team (e.g., E. Rykoff on 20241103) led to the following distinction of terminology that served to name the ITL memory effect in question. The main difference between “fluorescence” and “phosphorescence” is in that the former is considered prompt re-emission and the later could be re-emission following a finite characteristic time constant. Characteristic time constants are in the nanosecond scale for fluorescence, while for phosphorescence it would be in the milliseconds to seconds range. For the purpose of this discussion, we adopt the word “phosphorescence” to refer to the memory effect present in some ITL sensors.
6. Lesser mentioned that the wax-based resist **fluoresces** (that would be the prompt mechanism with very short relaxation time). If there is any such residual material between the coating and the passivated silicon, it would be natural to expect a halo that would accompany any sharp (PSF-scale features) illumination that passes through these “stains” on the sensor surface: a scatter term with low integrated amplitude, whose scale should depend upon the re-emission wavelength. This has not yet been seen in lab data but may appear once the Camera goes on-sky.

7.4.1 Measurement techniques for detecting and quantifying phosphorescence

We mentioned above that certain phosphorescent morphologies strongly resemble the “coffee stains” seen on the same (ITL) sensors. It would be logical to ask whether all “coffee stains” are accompanied by phosphorescence patterns. If this is the case, characterization and calibration of the deferred, phosphorescent signal could be extracted as a byproduct of the “cof-

fee stain” characterization⁶. The general answer to this question is that there is no simple relationship between the two, and the discussion that follows describes how this conclusion was established.

In this section, we describe the methods used to identify the transient term we consider phosphorescence in the ITL sensors, and we list the regions where it was detected. Following that, we describe in some detail the kinematics of its expression (cherry picking specific easy-to-measure cases), together with their wavelength and excitation flux-level dependencies.

We parasitically used a series of B-protocol persistence acquisitions executed for the purpose of tuning the operation of e2v sensors. The reason for this was that the ITL operating parameters were left unchanged from run to run, and thereby provided a large number of instances of identical EO measurement conditions (although the acquisitions were captured over a span of a few weeks). These EO runs acquired a series of dark images (with the nominal 15 s integration time, or ‘EXPTIME’) that followed a deliberate overexposure and readout of a FLAT (CCOB LED ‘red’, target signal 400 ke⁻/pix). The dark images acquired in succession following the FLAT image recorded the re-emitted or deferred signal collected within each 15 s period, and there were 20 such dark images acquired within each EO run. In all, we identified and analyzed a total of 22 runs that acquired this data, where the excitation flat properties were fixed. The first and the twentieth dark images were stacked and medianed following a nominal⁷ instrumental signal removal (ISR) step. The twentieth median dark images were then subtracted from the first median darks. This further suppressed any remaining ISR residuals from the resulting pixel data, which nominally now contain the *transient term* of the ITL phosphorescence: the component contained in the early (first) dark image (expressed in 15 s that has attenuated away by the final (twentieth) dark image, which was acquired 300 s (5 min) later.

⁶It is a great deal simpler to measure “coffee stains” with signal-to-noise limited only by illumination statistics, than it is to calibrate electron-level background structure emanating from exposure driven hysteresis.

⁷Two approaches for ISR were performed and the results were evaluated: *serial overscan only*, and *serial and parallel overscan joint correction*. With the downstream arithmetic step (twentieth median subtracted from the first median), the result is independent of any *master dark* subtraction utilized. The differences were very small, but consistent with the extra noise imparted by a superfluous, parallel overscan correction, which was further attenuated by the *median* operation on the image stack. We finalized these data products by using the *serial overscan only* ISR algorithm.

7.4.2 Results of phosphorescence detection in ITL sensors

Table 13 provides the EO run IDs analyzed according to the process outlined above. Figure 135 (top) displays the transient term in an 8×8 blocked image of the R00_SW1 sensor. (More examples can be found in Figs. 161 through 172 in Appendix C). These serve primarily to help identify which ITL sensors exhibit regions where we suspect presence of the phosphorescence effect. These full-resolution pixel images represent our most reliable data products for identifying individual pixels with significant transient signal.

A subset of the 88 ITL sensors, specifically those that show either high-signal diffuse, or morphologically unique structure in the transient term of the phosphorescence detected, are singled out to compare side-by-side with *blue* CCOB LED flat illumination. These are given in Appendix E, Figures 173 through 178. To illustrate their diagnostic power, Figure 134 provides an enlargement of four specific regions of sensor R02_S12 for better detail. It is apparent from these comparisons that generally, expression of phosphorescence has a complex relationship with the *much easier-to-detect* coffee stains (or other diffuse variations in quantum efficiency) seen on the same sensors: Presence of a coffee stain seen in flat field response may be suggestive of phosphorescence on the sensor, but accurate predictions are elusive. In some cases (as in Fig. 135 noted above), the phosphorescence appears to be correlated with the darker absorbed features of the coffee stain. In others (e.g., Figs. 134c & 134d), the opposite correlation is seen. In still other cases (e.g., Figs. 134e & 134f, Figs. 134g & 134h), there are regions of strong detail in the phosphorescence without very much coffee stain action at all. Our conclusions are that the presence of coffee stains does not provide a useful proxy for the phosphorescent properties of the sensor.

Table 13: Zephyr Scale E-numbers and corresponding SeqIDs analyzed to estimate phosphorescence in the 88 ITL sensors.

Run numbers and SeqIDs of first dark following trigger		
B-protocol runs, HVBias <i>off</i> , HVBias <i>on</i> for Corners		
E1003:20240920_000056	E1009:20240921_000222	E1003:20240920_000056
B-protocol runs, HVBias <i>on</i>		
Continued on next page		

Table 13 – continued from previous page

Run numbers and SeqIDs of first dark following trigger		
E1071:20240924_000300	E1110:20240926_000242	E1144:20240927_000369
E1146:20240928_001525	E1195:20241002_000235	E1245:20241003_000245
E1290:20241008_000286	E1329:20241011_001555	E1363:20241012_000546
E1392:20241014_000444	E1396:20241014_000701	E1411:20241015_000322
E1419:20241016_000397	E1429:20241016_000742	E1449:20241017_000548
E1497:20241020_000225	E1812:20241028_000481	E1880:20241030_000432
E2233:20241108_001468	E3380:20241130_000355	

While characterizing the phosphorescence expressed by ITL sensors using the data products described above, we have also identified correlations that concern the localized, phosphorescence centers that tend to appear as circular disks. While we typically see a dozen or so (on average) per sensor, those with larger amplitude are strongly associated with *vampire pixels* (which are easily identified by their localized flat field response). The correlation is not perfect, meaning that not all localized (circular) phosphorescence centers can be associated with *vampire pixels* but that nearly all *vampire pixels* express localized phosphorescence with some amplitude.

When data products of the 88 ITL sensors are inspected for transient phosphorescent response (cf. Table 14), very few, perhaps only a single sensor (R44_SW1), show insignificant phosphorescence. Although ~24% of the ITL sensors show diffuse phosphorescence, a majority of sensors (~83%) show spot-like phosphorescence centers without diffuse phosphorescence. Presence of diffuse phosphorescence probably can frustrate spot-like phosphorescence detection by eye, and the estimated frequency of the latter may serve as a lower limit to the true frequency. The case of R02_S12 (cf. Fig. 134) provides a good example of this.




Figure 134: Side-by-side comparisons of phosphorescence expression (left) vs. flat-field response to *blue* LED illumination (right) for 1500 pixel wide ROIs as indicated. Red markers indicate correspondence of isolated regions across images, while green markers indicate dust-related shadows and reflections seen only in flat response.




Figure 135: R00_SW1 image showing phosphorescence (top) with morphology similar to the “coffee stains” (bottom) observed with *blue* CCOB LED illumination. The phosphorescence acquired in dark exposures within the first 15 s following trigger (top) uses a logarithmic stretch with limits 5–25 e^-/pixel . The *blue* flat field (bottom) is displayed normalized, with 4% stretch limits (0.97 to 1.01), for a target signal level of $10^4 e^-/\text{pixel}$. Note that the phosphorescence pattern resembles the dark wisps in the flat (with opposite polarity) but that there are apparently no significant phosphorescence features corresponding to the bright wisps.

Table 14: Qualitative grouping of the 88 ITL sensors based on inspection of full resolution representations of Figures 161 through 172. In cases of spot-like phosphorescence, the numbers of features counted are given within parenthesis. Transient features appearing similar to *hot columns* or as other connected pixel groups are additionally signified with a double-plus (++).

Sensor Grouping		
Sensors exhibiting insignificant phosphorescence		
R44_SW1		
Spot-like phosphorescence (vampire transients)		
R00_SG0(>36)	R00_SG1(>36)	R00_SW0(>10)
R01_S00(>33)	R01_S01(>4)	R01_S02(>6)
R01_S10(>25)	R01_S11(18)	R01_S12(14)
R01_S20(>23)	R01_S21(>30)	R01_S22(>30)
R02_S00(>32++)	R02_S01(>36)	R02_S02(>28)
R02_S10(6)	R02_S11(>30)	R02_S12(>25)
R02_S20(>14)	R02_S21(>9)	R02_S22(>6++)
R03_S00(13)	R03_S01(12)	R03_S02(>19)
R03_S10(9)	R03_S11(3)	R03_S12(10)
R03_S20(9)	R03_S21(18++)	R03_S22(16)
R04_SG0(>12)	R04_SG1(>30++)	R04_SW0(25)
R04_SW1(>30)	R10_S00(>30)	R10_S01(9)
R10_S02(32)	R10_S11(16)	R10_S12(>26)
R10_S20(21)	R10_S21(>11++)	R10_S22(>10++)
R20_S00(2)	R20_S01(8)	R20_S02(7)
R20_S10(>35)	R20_S11(7)	R20_S12(5)
R20_S20(10)	R20_S21(5)	R20_S22(5)
R40_SG0(>50++)	R40_SG1(6++)	R40_SW0(6)
R40_SW1(8)	R41_S00(9++)	R41_S01(16)
R41_S02(10)	R41_S10(12)	R41_S11(3)
R41_S12(10++)	R41_S20(5++)	R41_S21(~30)
R41_S22(3)	R42_S00(24)	R42_S01(6)
R42_S02(>10)	R42_S10(4)	R42_S11(11)
R42_S12(33)	R42_S20(7)	R42_S21(5)
Continued on next page		

Table 14 – continued from previous page

Sensor Grouping		
R42_S22(4)	R43_S00(22++)	R43_S01(30)
R43_S02(19)	R43_S10(26)	R43_S12(8++)
R43_S21(14)	R43_S22(4)	R44_SG0(>12)
R44_SG1(>10)	R44_SW0(18)	
Segments exhibiting diffuse transient phosphorescence		
R00_SG1_C10-12,C03-05 (++)	R00_SW0_C17	R00_SW1_C** (++)
R01_S00_C13-14 (++)	R01_S01_C07,C16-17	R01_S10_C00-01,C14-16
R01_S20_C04-07	R01_S21_C06-07,C17	R01_S22_C00-01,C15-17
R02_S02_C03-04	R02_S11_C13-17,C07 (++)	R02_S12_C04-07,C10-12
R02_S20_C06-07	R04_SG1_C01,C11 (++)	R10_S10_C10,C16-17,C07
R40_SG0 (++)	R41_S21_C00,C10	R42_S00_C01,C07,C17
R43_S11 (++)	R43_S20_C00-01 (++)	R44_SG1_C07

The correspondence between *vampire pixels* and spot-like phosphorescence is laid out in Figure 136, for two prominent cases. These two *vampire pixels* may appear intrinsically different in that their flat-field responses **do** (or **do not**) exhibit a central bright pixel, which could aid in their identification. Details of the underlying distribution of trapped surface charges near the back-side electrode – or variations in the conductive properties of the same – apparently drive these details of the flat field response. However, it remains intriguing that these surface electrostatic properties are accompanied by an unmistakable transient phosphorescence signature.

A curious aspect of the phosphorescence seen in ITL sensors lies in its voltage (HV Bias) dependence. The HV Bias, when turned on, reduces lateral diffusion of the photo-conversions and thereby maintains PSF image quality. In Figure 137 we compare side-by-side several phosphorescent regions with both HV Bias states (off and on). There appears to be no trend that lends to predictability in these cases. In the cases of vampire pixels (R03_S10 & R20_S20), the geometry of the phosphorescence is indeed very sensitive to the HV Bias states (cf. Figs 137a vs. 137b; 137c vs. 137d). These might be understood qualitatively. However, for the diffuse




Figure 136: Vampire pixel comparisons between their flat field response and their transient phosphorescence. Signal levels are given (relative for flat field response, absolute electrons per 15 s following overexposure for transient phosphorescence). The relative flat field response amplitudes swing between 0.2 & 16 (reaching full well) for R03_S10, and between 0.4 & 8 for R20_S20. The transient phosphorescence response also reaches nominal full well (135 $\text{ke}^-/\text{pix}/15\text{s}$ for the central pixel) for R03_S10, and a lower amplitude (3–4 $\text{ke}^-/\text{pix}/15\text{s}$ for several hundred pixels) is reached for R20_S20.

phosphorescence examples, the expression appears to vanish entirely (R43_S11, Fig. 137e) or become significantly stronger, together with morphological changes (R43_S20, Fig. 137g) when the HV Bias is switched off.

7.4.3 Summary of phosphorescence in ITL sensors

From the investigations above, we find that a number of small but perhaps significant systematic effects appear to be joined by the presence of deferred signal which we loosely call *phosphorescence*. It is a hysteretic phenomenon and may be corrected pixel-by-pixel with adequate characterization. There appear to be multiple time constants in its decay (ranging from a short timescale of 11 s through longer timescales of 320 s) with roughly 75% of the integrated phosphorescence belonging to the longest timescales and only about 5% associated to the shortest timescales. This statement is based on a sample kinetics study performed on a specific ROI explored in Appendix F. We also find that the phosphorescent response to illumination is dependent on wavelength and illumination level (cf. Appendix G), which in certain cases appear to present a proportional response to trigger illumination.

In principle, each group of pixels (or each individual pixel) may be characterized by specific populations of characteristic phosphorescence kinetics as well as trigger exposure wavelength- and illumination-response. A dedicated set of measurements that acquire sequential dark images following a trigger exposure may be performed repeatedly until adequate statistics are achieved.

We consistently find phosphorescent counterparts to both *vampire pixels* and often (but not always) to *coffee stains* that appear in short wavelength flat field response. These facts provide strong clues that the phosphorescent signal is tied to surface detail of the backside window or equipotential. The fact that the time constants measured are comparable to those found for *persistence* in the other set of sensors (e2v) cast some doubt on the presumed mechanism and the name choice of the phenomenon. We do not know for certain that what we call *phosphorescence* isn't indeed a deferred signal caused by an temporarily trapped, then liberated charge, rather than a deferred photon emission as the adopted name would imply. We believe this suggestion is bolstered by the observation that the phosphorescent signal is strongly modified by the electrical state of the *HV Bias* switch to these sensors (cf. Fig. 137).

Whatever the underlying cause of the *phosphorescence* signal, it is a contribution to the instrumental signature in science data that attenuates over time, and is generally a hysteresis that




Figure 137: Comparisons of transient phosphorescence between conditions where HV Bias is off (left) vs. on (right). Four different ROIs are shown, but with image scales set to match across HV Bias conditions.

responds to wavelength, illumination history, and has strong position dependence to its expression with time constants that may also be position dependent. Whether it is corrected to first order as part of the Instrument Signature Removal (ISR) effort or not, sources extracted from regions suffering significant *phosphorescence* may be labeled as *suspect* and two-sided (or asymmetric errors) may apply to the variance plane when such sources are characterized or fitted.

Draft

7.5 Summary

Tree rings are concentric variations in silicon doping concentration observed in flat images. The centers of these tree rings have four possible positions relative to the sensor center due to the way CCDs are cut from wafers. Recent data with 0V back bias voltage and a diffuser installed has allowed for better measurement of individual tree ring centers. The radial study confirms that tree rings are not perfectly circular. The use of a diffuser reduces the impact of other patterns, such as CMB and weather patterns, improving the visibility of tree rings. Voltage dependency tests show that back bias voltage reduces the tree ring effect. The tree ring effect is consistent across different wavelengths, although other sensor effects dominate at lower and higher wavelengths.

ITL dips were observed in LSST ComCam on-sky data as bleed trails from bright stars with a dark core compared to the wings of the trail. Tests using spots and rectangles projected onto the focal plane did not find evidence of ITL dips in the LSSTCam CCDs. The contrast and brightness achieved in the lab may not be representative of on-sky measurements, and differences in sequencer files and operating temperatures could explain the effect.

Vampire pixels are a feature found on some ITL sensors, characterized by a group of pixels with higher photo-response surrounded by pixels with lower photo-response. These complexes can extend to tens of pixels in radius and are also seen in phosphorescence response. Vampire pixels are less common in the LSSTCam than in ComCam. Prominent examples in the LSSTCam include features on sensors R01_S00_C13-4, R03_S10_C15, and R20_S20_C13. These complexes can cause astrometric and shape transfer errors and background estimation or source confusion errors.

Phosphorescence was discovered during the persistence optimization process in Run 7. It affects a subset of ITL sensors and is morphologically distinct from the persistence seen in e2v sensors. The memory effect is reminiscent of “coffee stains” seen in flat field response but with opposite polarity. The attenuation timescale is similar to that of e2v sensors, suggesting a similar de-trapping mechanism. The origin of the stains may be related to raised spots on the sensors’ backside surfaces. Phosphorescence is detected using a series of dark images following overexposure to a flat field.

The results show that phosphorescence is present in a subset of ITL sensors, with some sensors showing high-signal diffuse or morphologically unique structures. The presence of coffee

stains in flat field response may suggest phosphorescence, but the correlation is not consistent. Localized phosphorescence centers are often associated with vampire pixels. The voltage dependency of phosphorescence shows that HV Bias affects the geometry and intensity of the phosphorescence.

Draft

8 Overall Operation and Issues

In this section, we briefly report LSSTCam operational issues that were encountered in the Run 7 period.

8.1 Timeline of events and issues

The Gantt Chart below shows the timeline of event, not only limits to the actual EO testing period, but also includes the preparation period prior to the run. For details on individual runs, see Section 9.3 or visit Zephyr scale (rubinobs.atlassian.net access required). The blue bars show the actions that were taken while the red bars show the issues encountered. The issues corresponding to each red bar (and one of the milestones) are discussed in the indicated subsections.




Figure 138 shows the temperature trending of the Cryo plate, one of CCDs, R01_S10, and the REBs over the whole run period. These show smooth behavior except intervals where we had troubles of one kind or another.




Figure 138: Temperatures of Cryo plate, one of CCDs, R01_S10, and the REBs over the whole run period.

8.2 Camera control network performance

The Camera Control System (CCS) is a distributed system composed of processes running on servers in the summit computer room, desktop machines in the control room and the clean-room, as well as embedded computers (Hardware Control Units) in the camera utility trunk, and attached to the camera refrigeration system. The communication mechanism relies on a peer-to-peer communication system implemented using JGroups. This system functioned reliably with the LSSTCam at SLAC during the Run 6 testing, as well as with AuxTel and ComCam in Chile, but with the LSSTCam at Cerro Pachón we encountered previously unseen problems, the symptoms of which were apparently spontaneous loss of communication between camera subsystems (“CCS meltdowns”) which typically required manually shutting down a large fraction of the system to restore communication, rendering the setup unsustainable for long-term operations.

At SLAC the entire network was configured using a single network switch in the IR2 clean-room, but the setup in Chile was significantly more complex with different network switches on different levels of the observatory. We worked closely with the Rubin Chilean IT team to understand the network configuration, and to track down the cause of the problem. Key observations indicated that while normal CCS traffic was light (3.5 Mb/s), meltdowns generated traffic exceeding 2 Gb/s. Tests revealed that moderate multicast traffic (100 MB/s), deliberately created using iperf, could consistently induce the failures.

Ultimately a number of changes were made:

- The network switch controlling the White Room was replaced, and the network configuration in general was simplified to reduce the number of different switches used by LSSTCam.
- The total volume of network traffic generated by the camera was reduced and made more uniform over time by adjusting when large-volume telemetry, in particular from the focal plane and REB power systems, was published.
- Extensive diagnostics were added to CCS to help us identify the cause of the meltdowns. This was ultimately understood to be ‘retransmission storms’ triggered by the mechanism built-in to JGroups to recover from packet loss, but which was in fact causing far more message retransmissions than the network could handle. This was perhaps exacerbated by the fact that we were using a mixture of 10 Gb and 1 Gb connections, so

the control room servers (with 10 Gb connections) were able to generate far more traffic than the MOXA switch in LSSTCam (1 Gb) could handle.

- After discussing with the jGroups maintainers we upgraded the entire CCS from jGroups4 to jGroups5, and added some custom code to reduce the maximum volume of retransmission messages.
- We discovered that the Beckhoff PTP module used in the LSSTCam shutter (lsstcam-shutter01-ptp) was only capable of handling 100 Mb/s traffic, and this seemed to contribute to some of the problems. This was fixed by enabling IGMP in the MOXA network switch in the utility trunk to partition the network so that the CCS multicast traffic would not be sent to the PTP module.

Once these changes were made the network worked much more reliably, although there was one reoccurrence of the problem towards the end of the test period, caused by operating the camera in Guider mode with 400×400 ROIs, which at the current time has not been understood. While we do not think a single fix solved the problems, we are now in a much better position to diagnose and fix any future problems that arise.

8.3 REB PS power trip

In the evening of August 29, 2024, at 20:00 local time (00:00 UTC), two power units, R43 and R33, on RebPS P00, lost power. At that time, no one was at the summit.

After 20:38 local time, two people on shift went up to the White Room from the hotel and conducted an inspection, but they found nothing unusual.

When they arrived at the summit is unclear. Another incident occurred. R21/Reb[12] and R42 on P04 also lost power.

We suspect that this was caused by static electricity events. The Utility Trunk door appeared to be floating electrically, and the people who had worked on the LSSTCam in the White Room up to that date did not wear ESD straps. To address this issue, we grounded the Utility Trunk door and implemented an administrative control to ensure that everyone working around LSSTCam wore ESD straps. Since then, no REB power supplies have tripped off. A full report is available.

8.4 FES latch sensor failure during motion

At its arrival at the summit, the Filter Exchange System (FES) in LSSTCam was equipped with the new spare Autochanger (AC2, delivered to SLAC from France in fall 2023). AC2 was tested in LSSTCam in February 2024, just before the camera was shipped to Chile. For the first light, the plan is to use the well-known AC1, which, after four years of operation, underwent major maintenance in February 2024. AC2 and AC1 were swapped in LSSTCam in June 2024 at the summit, and Run 7 provided an opportunity to re-qualify AC1 following its maintenance.

Unfortunately, two weeks into Run 7, one PLC signal for AC1 (the X-open latch signal, related to the filter handling by the AC) was lost. Although the signal briefly returned, for safety reasons, the FES was not used until a dedicated intervention in November 2024.

During the maintenance of AC1 in February 2024, the cable associated with this PLC signal had been identified as a potential issue: the protective sheath around two cables associated with the AC latches (which secure the filter in the AC trucks) was identified as fragile and a potential dust source in LSSTCam. These cables, which change shape during filter movement, could generate dust over time. However, due to time constraints, the cables were not replaced during the February 2024 maintenance. Instead, an additional protective sheath was added to contain any dust. Unfortunately, the increased rigidity of the cables caused excessive strain on the connectors, leading to problems with both latches.

In November 2024, AC1 was removed from the camera, allowing the replacement of one of the two problematic cables: the latch X- probes cable, which was responsible for the faulty PLC signal causing the issue in September 2024. This repair resolved the problem. Due to tight scheduling, the second cable could not be replaced before AC1 was reinstalled in the camera. This second cable also showed an issue, preventing the full opening of the X+ latch. However all potential issues with this cable were resolved with AC1 installed in the camera by reducing the strain caused by the stiff cable sheath.

With these fixes, the “stiff cable issue” is now resolved. AC1 has been re-qualified through successful operation at the end of Run 7.

The full report is available.

8.5 PCS degradation

After a week of operation of the Pumped Coolant System (PCS), it was observed that its performance declined over time (Fig. 139). The PCS automatically shut down at 14:03 local time on October 15, 2024, resulting in LSSTCam running without cold cooling for 10 min. The REB temperatures reached nearly the maximum limit that triggers self-shutdown of the REBs, but the Camera was able to continue operating as it regained cooling capacity when the PCS was powered on again after 10 min.




Figure 139: The trend of TXV Bulb temperature at PCS Stage 1. The increasing trend indicated the degradation over time.

PCS performance degradation was observed following the event. Several mitigation strategies were used as this issue continued, including:

1. Closing the Glycol valve for the PCS for 70 s (subsequently increased to 120 s).
2. Stopping the refrigeration compressor while allowing refrigerant flow to continue.
3. Adding 250 g refrigerant, or more.

However, on November 10, the chiller no longer provided sufficient cooling capacity, resulting in the Camera being placed in a degraded operational mode with inadequate cold cooling while cryo cooling remained sufficient.

A concern was raised that the PSC behavior resulted from the loss of refrigerant. An action was taken to recover the PCS refrigerant for weighing purposes. The measured weight was 6.404 kg, which deviated from the expected value of 6.804 kg. The engineer deemed this discrepancy insufficient for the degradation of the system's performance. The original mass of

refrigerant (6.804 kg) was returned to the system. Subsequently, the superheat valve of the TXV was adjusted to fully open the valve. These actions resulted in a week of stable operation at the end of Run 7.

The definitive resolution has not yet been made. Further analysis is required.

Draft

8.6 R24/Reb0 and UT leak fault issue

We were recovering operations from the PCS issue on November 10, 2024, at 10:17 UTC when the PCS could not continue to operate. After the PCS tripped off, we let it sit for several hours. We warmed up the cold plate to -10 C.

After a few more hours, we attempted to reestablish the cold condition. The PCS set point was lowered gradually, and we were adding heat load by powering on REBs. However, during this process, R24/Reb0 tripped off due to an error message: "Execution of command 'powerRebOn' failed unexpectedly due to: REB 0 power on sequencing failed: 0x00080020 (OD voltage or current out of range)." We power cycled the RebPS P02, which powers R24/Reb0, and R24/Reb0 tripped off again with the same error.

A few minutes later, UTLeakFault became asserted, and the Camera was fully shut down. A crew went to the White Room but did not find any water leaks, nor did the leak detector. We have not identified the cause. (We note that a few days later, we found a few drops of Dynalene in the UT, although it was found that it was not enough to trigger the UT leak detector. The source of the very small leak was sealed with TorrSeal.)

On November 15, 2024, we set the Dynalene temperature to 4 C, then changed it to 8 C and 11 C. The temperature at the Camera side reached 12.5 C. RebPS/BoardTemp6 reached 13 C. Under these conditions, R24/Reb0 came on without any issues. Therefore, we determined that the low RebPS/PS02/BoardTemp6 of around 9 C was the cause of the problem. We note that this temperature is not a strict limit. Some of the power supplies could run at this temperature. Also once it powers up the board temperature (or Dynalene temperature) needs to be as low as 9 C (4 C) so not to trip off the power supply by hitting the high temperature limit (48 C) when the HV unit is on.

Regarding the UT Leak Fault issue, we later realized that it is crucial to distinguish between "UT Leak" and "UT Leak Fault." "UT Leak" refers to an actual leak detected, while "UT Leak Fault" indicates an issue with the leak detector itself. Subsequent investigations suggested that the DC-DC converter supplying power to the UT leak detector may not have been adequately loaded, leading to its shutdown and triggering the "UT Leak Fault." We added a load resistor to the DC-DC converter to prevent future trouble.

A detailed note is available.

8.7 Data corruption

After the temporary fix of the PCS degradation, when we tried to turn off the CCDs, we observed horizontal stripes in the raw images (Fig. 141), which made the overscan subtracted-image looks not normal (Figure 140).




Figure 140: The first observation of the data corruption in the full focal plane image.




Figure 141: A close-up view of the section of the image showing the horizontal stripes (without the overscan subtraction).

We rebooted the DAQ, and power cycled the CCDs, REBs, focal-plane subsystem and image-handling subsystem. However, none of these actions fixed this issue.

Restarting Data Store Reconfigurable Computing Elements (RCEs finally brought back the system normal.

FRACAS ticket 256 was filed.

8.8 Guider high gain issue

The guider sensors encountered issues while operating in guider mode, failing to produce reasonable images. The bias level was abnormally low, measured at 1.0--1.5k instead of the expected 15k ADU. Additionally, the CCOB flashes were not recorded correctly. Although two flashes of 1.2s each occurred, the sensor did not provide a meaningful response. At the time, the guider was configured with an integration time of 50 ms and a readout time of approximately 70 ms. The sensor should have displayed flat illumination across several frames, but the mean image value remained unchanged, with only minimal variations in the standard deviation and maximum values.

During troubleshooting, discussions were held on Slack, and diagnostic steps were undertaken. These conversations highlighted the behavior of the sensor, with reference to previous image acquisitions. For example, on September 18, 2024 the guider had successfully produced a mean count of 15k during image acquisition with the partial focal plane, meeting expectations. Furthermore, the guider appeared to function normally in science mode, where no abnormalities were observed.

Immediate actions included performing a full power cycle and switching back to science mode. However, the guider tripped off during power-on, requiring another power cycle. The issue was identified as related to REB PS 4, which had tripped with an error code indicating the clock high voltage was too low (0x00280040). This necessitated a reset of the RCE following the power cycling of the CCDs/REBs. Despite these challenges, the process of switching to guider mode was repeated, and this time the system returned to normal. The baseline signal level stabilized at 1.5k, and the sensor responded correctly to flashes.

Although the immediate issue was resolved, no definitive root cause was identified. Further investigation is needed to understand the conditions that led to the tripping of REB PS 4 and to prevent such an occurrence in the future.

FRACAS ticket 246 was filed.

9 Conclusions

9.1 Run 7 final operating parameters

This section describes the conclusions of Run 7 optimization and the operating conditions of the camera. Decisions regarding these parameters were based upon the results of the 5.1voltage optimization, 5.2sequencer optimization, and 5.2.3thermal optimization.

9.1.1 Voltage conditions

Table 15: Voltage Conditions. ITL voltages have not changed since Run 5, and are reported here.

Parameter	Run 5 (ITL)	Run 5 (dp93; e2v)	Run 7 (dp80; new voltage for e2v)
pclkHigh	2.0	3.3	2.0
pclkLow	−8.0	−6.0	−6.0
dpclk	10.0	9.3	8.0
sclkHigh	5.0	3.9	3.55
sclkLow	−5.0	−5.4	−5.75
rgHigh	8.0	6.1	5.01
rgLow	−2.0	−4.0	−4.99
rd	13.0	11.6	10.5
od	26.9	23.4	22.3
og	−2.0	−3.4	−3.75
gd	20.0	26.0	26.0

9.1.2 Sequencer conditions

Table 16: Sequencer conditions

Detector type	File name
e2v	FP_E2V_2s_l3cp_v30.seq
ITL	FP_ITL_2s_l3cp_v30.seq

- v30 sequencers are identical to the FP_ITL_2s_l3cp_v29_Noppp.seq and FP_E2V_2s_l3cp_v29_NopSf.seq

All sequencer files can be found in the GitHub repository.

For both sequencers, idle flush is disabled

9.2 Record runs

This section describes the record runs for Run 7, i.e., the electro-optical testing runs from which the Run 7 performance characterization described in this note was derived. All runs use our camera operating configuration, unless otherwise noted.

Run Type	Run ID	Notes
B protocol	E1880	Initial B protocol run taken with the new v30 definition.
	E2233	dp80, first run after full CCS system reboot
	E3380	First B protocol post-chiller recovery
PTC	E3630	Low flux red LED PTC, ND1 filter installed. Final operating conditions of camera.
	E3577	Dense nm960 PTC. Final operating conditions of camera.
	E2237	Final operating conditions of camera. Red LED dense. Acquired after CCS subsystem reboot.
	E748	Final operating conditions of camera. nm960 dense
	E2016	Final operating conditions of camera. Super dense red LED. HV Bias off for R13/Reb2. jGroups meltdown interrupted acquisitions, restarted
	E1886	Final operating conditions of camera. Red LED dense. Dark interleaving between flat pairs
	E1881	Final operating conditions of camera. Red LED dense. No dark interleaving between flat pairs
Gain Stability	E1955	6h Stability run 10k 750 nm V30, dp80, idle flush disabled
	E2008	6h Stability run 10k 750 nm V30, dp80, idle flush disabled, after zero-ing CCOB

Run Type	Run ID	Notes
	E2136	15 s darks with variable delays between acquisitions
	E2236	Dark and biases with consistent delays between acquisitions
	E2330	Dark images with delay between acquisitions.
Long dark acquisitions	E3540	900 s dark. Shutter closed.
	E3539	900 s dark. Shutter closed.
	E3538	900 s dark. Shutter opened.
Projector acquisitions	E2184	10 30 s dark images to capture background pattern. E2V:v29Nop, ITL:v29Nopp
	E2181	Flat pairs from 2–60 s in 2 s intervals. Two 15 s darks interleaved after flat acquisition. Rectangle on C10 amplifier. E2V:v29Nop, ITL:v29Nopp
OpSim runs	E2330	Short dark sequence, filter changes in headers through OCS
	E2328	Flats with shutter-controlled exposure
Phosphorescence datasets	E2015	10 flats at 10 ke ⁻ followed by 10×15 s darks
	E2014	1 flat at 10 ke ⁻ followed by 10×15 s darks
	E2013	10 flats at 10 ke ⁻ followed by 10×15 s darks. Interleaved biases with the darks
	E2012	10 flats at 1 ke ⁻ followed by 10×15 s darks
	E2011	20 flats at 10 ke ⁻ followed by 10×15 s darks

9.3 Other runs of relevance

Runs that use the Run 7 final camera operating configuration (Sec. 9.1) are denoted with **bold run ID**.

Table 18: B Protocol Runs

Run ID	Notes
E3380	First B protocol post-chiller recovery. v30, dp80, idle flush disabled.

Run ID	Notes
E2233	Identical to E1880. Acquired after CCS subsystem reboot. dp80, idle flush disabled.
E1880	Camera operating configuration
E1812	v29 NopSf (no pocket serial flush running for both e2v and ITL clear sequencers). dp80 voltages, idle flush ?? [likely disabled but verification needed]
E1497	v29 Nop sequencer, dp80, idle flush ?? [likely disabled but verification needed]
E1429	First dp84 run. v29, idle flush disabled
E1419	First dp88 run. v29, idle flush disabled
E1411	First dp865 run. v29, idle flush disabled
E1396	First dp80 run. v29 nonoverlapping sequencer, idle flush enabled
E1392	First dp80 run. v29 sequencer, idle flush enabled
E1290	using Guide sensors as guiders. v29, dp93, idle flush enabled
E1245	Refrigeration system software update mid-run. v29 halfoverlapping sequencer. dp93, idle flush enabled
E1195	v29 overlap113 sequencer (5% overlap). dp93, idle flush enabled
E1146	First run with v29 nonoverlapping. dp93, idle flush enabled
E1144	First run with v29 Nop. dp93, idle flush enabled
E1110	v29 run. dp93, idle flush enabled
E1071	SOURCE = 63 in calib3.cfg. First run with HV on. dp93, v26 sequencer, idle flush enabled

Table 19: PTC Runs

Run ID	Notes
E3630	Low flux red LED PTC, ND1 filter installed. Final operating conditions of camera.
E3577	Dense nm960 PTC. Final operating conditions of camera.
E2237	Final operating conditions of camera. Red LED dense. Acquired after CCS subsystem reboot.
E748	Final operating conditions of camera. nm960 dense
E2016	Final operating conditions of camera. Super dense red LED. HV Bias off for R13/Reb2. jGroups meltdown interrupted acquisitions, restarted
E1886	Final operating conditions of camera. Red LED dense. Dark interleaving between flat pairs
E1881	Final operating conditions of camera. Red LED dense. No dark interleaving between flat pairs
E1765	Dense PTC, red, thresholded dark interleaves, overlaps in signal level for adjacent LED currents. v29 Nop sequencer, idle flush state unknown
E1495	dp80, nopp config. Idle flush ??
E1364	v29, dp80, idle flush state unknown. Possible incomplete data transfer
E1335	dp80 configuration, v29, idle flush ??.
E1275	Ordered flats. Failed dark interleaving, incomplete data transfer. v29 sequencer.
E1259	Randomized flats. v29 sequencer.
E1258	Randomized flux levels. Starting with 3 preimages, then 100 15 s darks, then PTC set. No dark interleaving. v29 sequencer.
E1247	Re-do of E1188 (which lacked PD data). v29HalfOverlapping., Added pre-image acquisition to PTC-Red cfg file.

Run ID	Notes
E1212	5% overlapping sequencer
E1145	No pocket sequencer
E1113	v29 sequencer
E749	v26, dp93, idle flush enabled. First PTC of run.

Table 20: Long Dark Acquisitions

Run ID	Notes
E3540	900 s dark. Shutter closed.
E3539	900 s dark. Shutter closed.
E3538	900 s dark. Shutter opened.
E1140	Empty frame filter, shutter open, 24 V clean and dirty FES changer powered off, one 900 s dark image only.
E1117	900 s dark. r filter, shutter open.
E1116	900 s dark. y filter, shutter open.
E1115	900 s dark. g filter, shutter open.
E1114	900 s dark. EF filter, shutter open.
E1076	PH filter in place. Shutter open. v26 no RG
E1075	PH filter in place. v26 no RG

Table 21: Projector Acquisitions

Run ID	Notes
E2184	10 30 s dark images to capture background pattern
E2181	Flat pairs from 2–60 s in 2 s intervals. Two 15 s darks interleaved after flat acquisition. Rectangle on C10 amplifier. e2v:v29Nop, ITL:v29Nopp
E1586	One 100 s flat exposure, spots moved to selected phosphorescent regions. e2v:v29Nop, ITL:v29Nopp
E1558	Flat pairs, fine scan in flux from 1–100 s in 1 s intervals. e2v:v29Nop, ITL:v29Nopp

Run ID	Notes
E1553	Flat pairs, coarse scan in flux from 5–120 s in 5 s intervals. e2v:v29Nop, ITL:v29Nopp

Table 22: OpSim Runs

Run ID	Notes
E3629	Mock OCS calibrations, failed
E3576	Mock OCS calibrations, failed
E3570	Mock OCS calibrations, failed
E2330	Short dark sequence, filter changes in headers through OCS
E2329	Mock OCS calibrations, failed
E2328	Flats with shutter-controlled exposure
E2283	Full night of OpSim flats, failed
E2280	Mock OCS calibrations, failed
E2279	Mock OCS calibrations, failed
E1717	Long dark sequence, no filter changes
E1657	10 hour OpSim dark run, ~50% of darks were acquired properly
E1414	30 minutes OpSim run with shutter control, filter change, and realistic survey cadence
E1403	30 minutes OpSim run with shutter control, filter change, and realistic survey cadence
E1255	30 minutes OpSim run with shutter control, filter change, and realistic survey cadence
E1254	30 minutes OpSim run with shutter control, filter change, and realistic survey cadence
E1092	30 minutes OpSim run with shutter control, filter change, and realistic survey cadence

Table 23: Phosphorescence Datasets

Run ID	Notes
E2015	10 flats at 10 ke ⁻ followed by 10×15 s darks

Run ID	Notes
E2014	1 flat at 10 ke ⁻ followed by 10×15 s darks
E2013	10 flats at 10 ke ⁻ followed by 10×15 s darks. Interleaved biases with the darks
E2012	10 flats at 1 ke ⁻ followed by 10×15 s darks
E2011	20 flats at 10 ke ⁻ followed by 10×15 s darks

Table 24: Tree Ring Flats

Run ID	Notes
E1050	Red LED. HV off. Diffuser installed.
E1052	Blue LED. HV off. Diffuser installed.
E1053	Nm750 LED. HV off. Diffuser installed.
E1055	Nm850 LED. HV off. Diffuser installed.
E1056	Nm960 LED. HV off. Diffuser installed.
E1021	Red LED. HV off. Diffuser removed.
E1023	Blue LED. HV off. Diffuser removed.
E1024	Nm750 LED. HV off. Diffuser removed.
E1025	Nm850 LED. HV off. Diffuser removed.
E1026	Nm960 LED. HV off. Diffuser removed.

Table 25: Gain Stability Runs

Run ID	Notes
E1955	6h Stability run 10k 750 nm V30, dp80, idle flush disabled
E2008	6h Stability run 10k 750 nm V30, dp80, idle flush disabled, after zero-ing CCOB
E1968	6h Stability run 2k 750 nm V30, dp80, idle flush disabled
E1367	Changing PCS setpoint mid run., PCS changed from -45 deg C to -47 deg C at 10:40:06 AM UTC. 6h, 50k at 750 nm, v29 seq, dp80 config.
E1362	dp80, partial data ingestion. v29 sequencer. 6h 10k at 750 nm.

Run ID	Notes
E756	dp 80, v29 sequencer. 6h 10k at 750 nm. Partial data ingestion.
E1496	dp80, nopp config, 12h 750 nm at 10k

Table 26: Persistence Datasets

Run ID	Notes
E2286	30k uv flash with increased hilim
E1507	dp80, uv led @ 30k
E1506	dp80, uv led @ 10k
E1505	dp80, uv led @ 5k
E1504	dp80, uv led @ 3k
E1503	dp80, uv led @ 1k
E1502	dp80, blue led @ 1k
E1501	dp80, blue led @ 3k
E1500	dp80, blue led @ 5k
E1499	dp80, blue led @ 10k
E1498	dp80, blue led @ 30k
E1494	dp80, nm960 led @ 50k
E1493	dp80, nm850 led @ 50k
E1492	dp80, nm750 led @ 50k
E1491	dp80, blue led @ 50k
E1490	dp80, red led @ 50k
E1489	dp80, nm960 led @ 150k
E1488	dp80, nm850 led @ 150k
E1487	dp80, nm750 led @ 150k
E1486	dp80, red led @ 150k
E1485	dp80, blue led @ 150k
E1484	dp80, blue led @ 400k
E1483	dp80, red led @ 400k
E1479	dp80, nm750 led @ 400k
E1478	dp80, nm960 led @ 400k
E1477	dp80, nm850 led @ 400k

Table 27: Guider ROI Acquisitions

Run ID	Notes
E1509	ROI reference dataset
E1510	ROI crossing amplifier segments
E1518	200ms integration time
E1519	100ms integration time
E1508	50ms integration time
E1520	400x400 pixel ROIs
E1511	200x200 pixel ROIs
E1521	100x100 pixel ROIs
E1512	New row from reference dataset
E1513	New column from reference dataset
E1514	New column and row from reference dataset
E1517	Different row for sensors on the same REB

A Reference figures

A.1 From Sections 4 and 6.1

A.1.1 Focal-plane measurement figures

One of two common plots from Sections 4 and 6.1 is a focal-plane measurement plot, showing eo-pipe measurements in different runs. These plots are constructed to show measurements arranged localized to the LSSTCam focal plane orientation. Different sensors are plotted in different colors, with one point plotted for each amplifier on each sensor. The plots are arranged such that an apples-to-apples comparison is made between two different runs (see Fig. 142). On one axis is the measurements from one run, while the other axis has the measurements from a different run. If the measurements are consistent, the points should fall along the identity line. Deviations from the identity line indicate that one run is measuring the quantity at a higher or lower value than the other run. This plot is useful for identifying localized deviations between different runs, and also e2v and ITL specific behavior.




Figure 142: Full focal plane measurements of divisadero tearing between runs 13550 and E1071, with populations separated by manufacturer type.

A.1.2 Differential histograms

Another common plot in Sections 4 and 6.1 is a histogram, where the histogram bins tally differences between measurements of a given quantity for each amplifier. These histograms are commonly generated by detector type, to show any type-dependent differences. In the case of no significant differences, these histograms should be Gaussians centered around zero, with no biases in one direction or another.




Figure 143: Differential histogram of divisadero tearing between runs 13550 and E1071, with populations separated by manufacturer type.

A.2 Web report reference figures

A tool for rapidly reviewing LSSTCam EO data is the web reports, which provide full focal plane, raft-level, and sensor-level figures of reference for prompt interpretation of analyses. Web reports from Runs 6 and 7 are available.

A.2.1 Focal plane level

A.2.1.1 Focal plane mosaics for different quantities Focal plane mosaics represent amplifier-level EO test quantities arrayed as fully assembled focal planes, with amplifiers colored according to the associated measurement. The full focal plane mosaics are generated for most eo-pipe parameters. These visualizations show differences between e2v and ITL sensor performance across the focal plane, as well as variations on the sensor and amplifier levels (Fig. 144).




Figure 144: A focal plane mosaic from run E1881 for PTC gain. Depending on the operating mode of the wavefront and guider sensors, images from these sensors are stored differently. As a result, some EO quantities are not computed, and therefore not shown in the mosaics.

A.2.1.2 Histograms In addition to focal-plane mosaics, histograms are created for all eo-pipe metrics. These histograms have bin widths and ranges specific to each parameter. For some parameters, the histograms provide convenient visualizations to quantify differences between e2v and ITL performance, such as Figure 145 for PTC a_{00} .




Figure 145: A histogram from run E1881 for PTC a_{00}

A.2.2 Raft level

A.2.2.1 Correlation figures Correlation figures are created on the raft level for all science rafts. Two types of correlation figures are created; one for the imaging region, and one for the overscan region. The correlation computed is a Pearson correlation coefficient (Fig. 146).




Figure 146: Overscan correlations for R22_S11 in run E1880.

A.2.2.2 Lambda mosaics As a part of the standard B protocol acquisitions, flats are taken in different LEDs to provide a chromatic response across the LSST bandpass. The mosaic is assembled on the raft level, and short-wavelength flats show the laser annealing pattern characteristic to e2v sensors (Fig. 147).




Figure 147: Blue LED mosaic for R22 from run E1880.

A.2.2.3 Calibration frames For runs with bias, dark, and flat frame acquisitions (most commonly B protocols), combined calibration bias and dark frames are produced and assembled on the raft level. These mosaics are gain corrected (see Fig. 148 for an example of a calibration dark).




Figure 148: Calibration dark for R22 from run E1880, extracted using B protocol dark sequence.

A.2.2.4 Bias stability Bias stability figures are created for acquisition sequences that acquire multiple bias frames. Three different bias stability plots are created: amplifier-wise mean vs. time, amplifier-wise standard deviation vs. time, amplifier-wise mean vs. time for a region covering the readout corner. An example of the latter is shown in Figure 149.

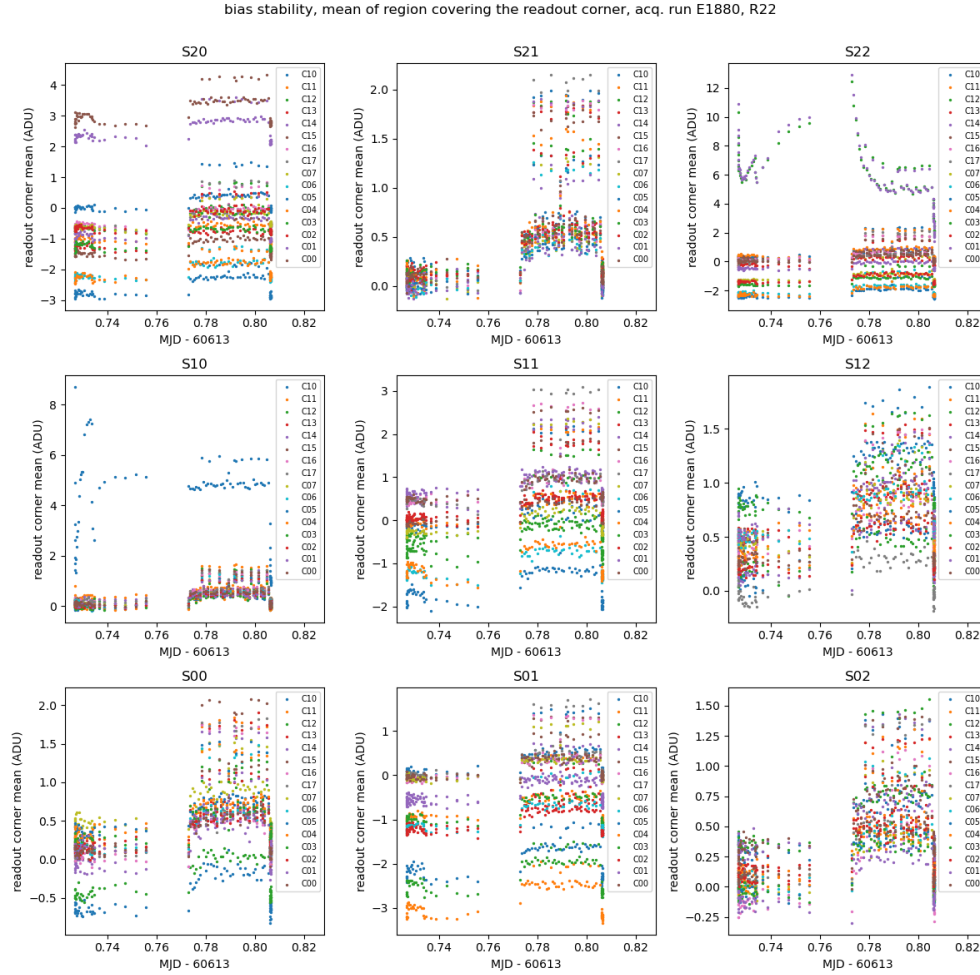


Figure 149: Bias stability mean vs. time for readout corners of the amplifiers organized by sensor in R22, from run E1880.

A.2.2.5 Divisadero profiles Divisadero tearing profiles are created for each raft, with divisadero response grouped along the mid-line break. Dashed lines indicate regions where divisadero effects on the response are expected (Fig. 150).

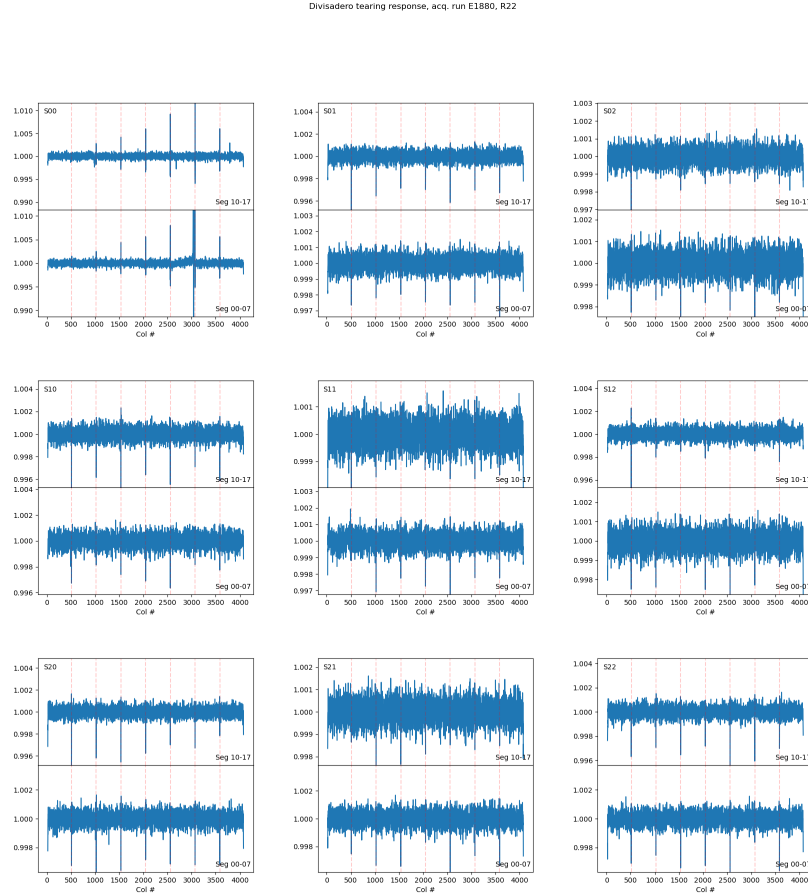


Figure 150: Divisadero tearing profile for R22 from run E1880.

A.2.3 Sensor level

A.2.3.1 Bias profiles Bias profile plots are created for each sensor, along both serial and parallel directions. Amplifiers are plotted separately to allow amplifier-dependent response to be identified. Different columns/rows are plotted in different colors for identification of problematic columns/rows (Fig. 151).

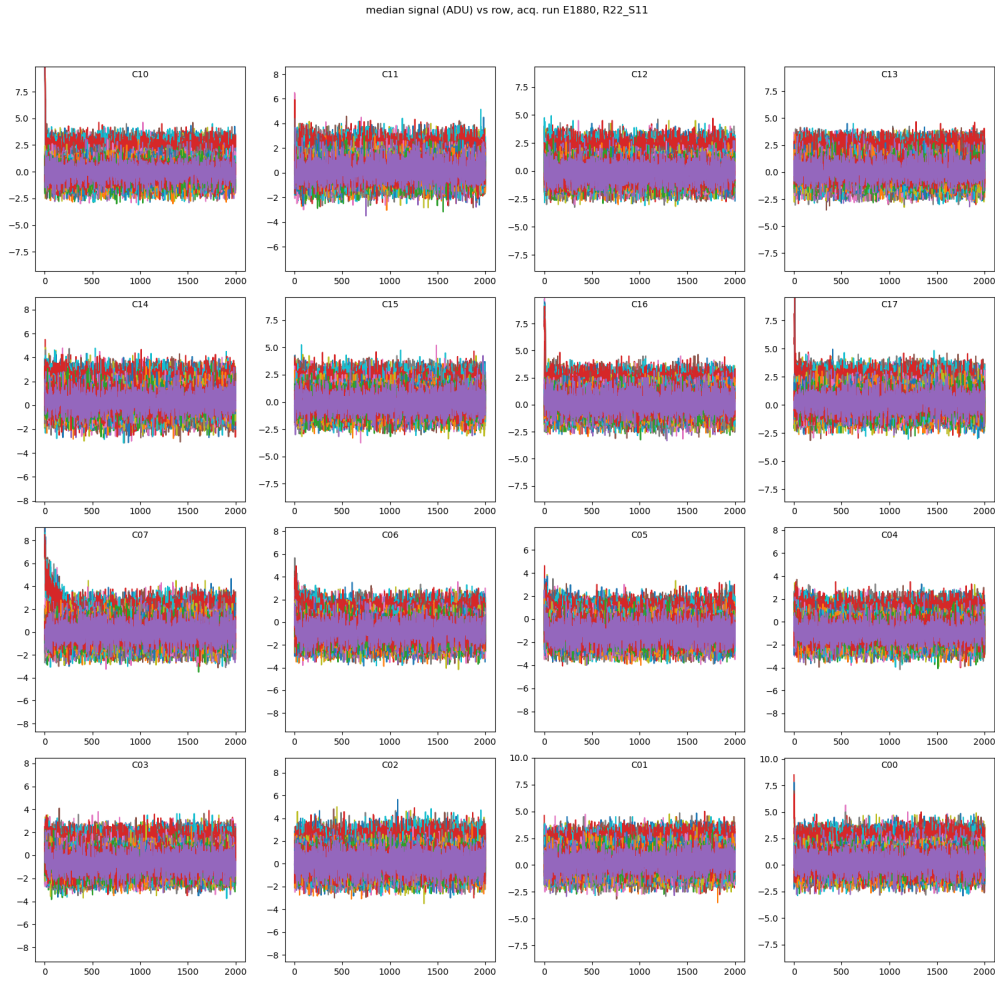


Figure 151: Median bias profile in the parallel direction for R22_S11 from run E1880.

A.2.3.2 PTCs PTC plots are created for each sensor, and separated by amplifier. For each amplifier, a PTC curve is created. PTC gain, a_{00} , and PTC turnoff (in ADU) is shown for each amplifier. The red points in PTC curves (Fig. 152) denote all measurements before PTC turnoff.

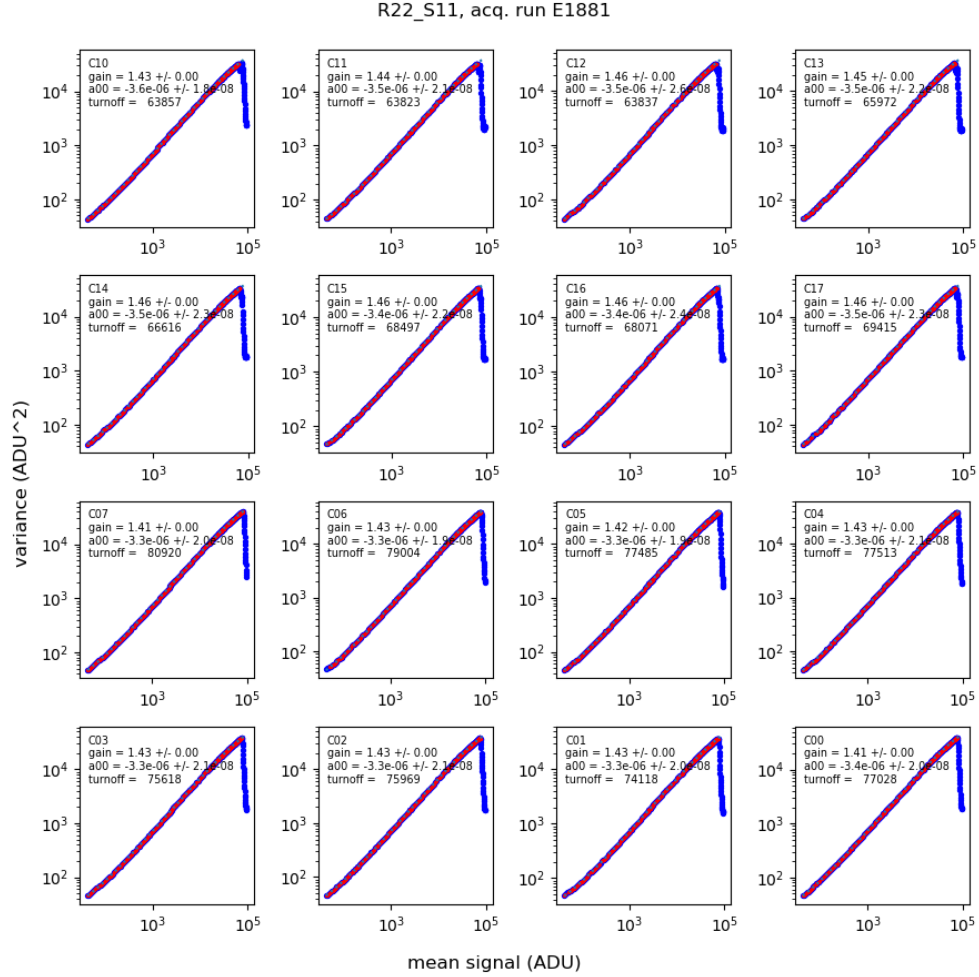


Figure 152: PTC plots for R22_S11 from run E1881.

A.2.3.3 Nonlinearity Linearity plots are created for each sensor, and separated by amplifier. For each amplifier, a linearity curve is created showing photodiode current integral vs. e-/pixel for all flat pairs. Maximum fractional deviation, maximum observed signal, and linearity turnoff (in ADU) are shown for each amplifier. The red points in the linearity curves (Fig. 153) denote all measurements between 100 e-/pixel and linearity turnoff.

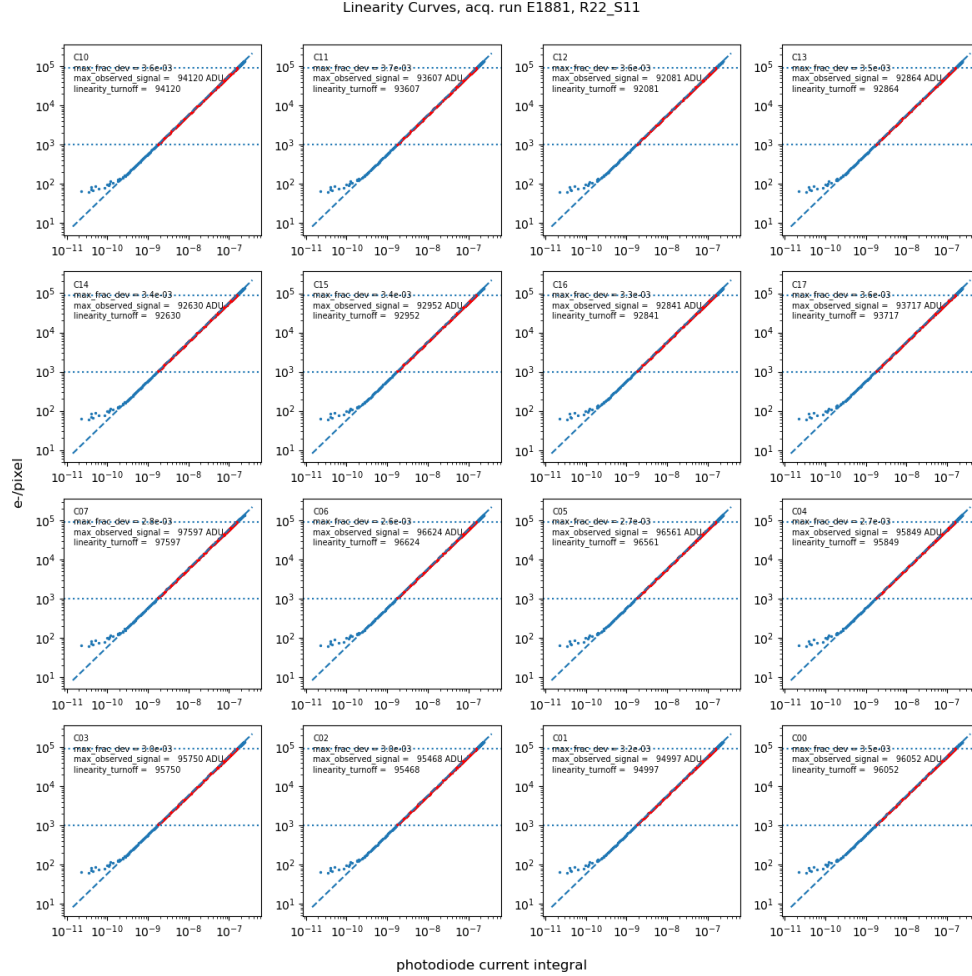


Figure 153: Linearity plots for R22_S11 from run E1881.

Plots of linearity residuals are also generated, showing the fractional residuals for each amplifier as a function of the photodiode current integral (Fig. 154).

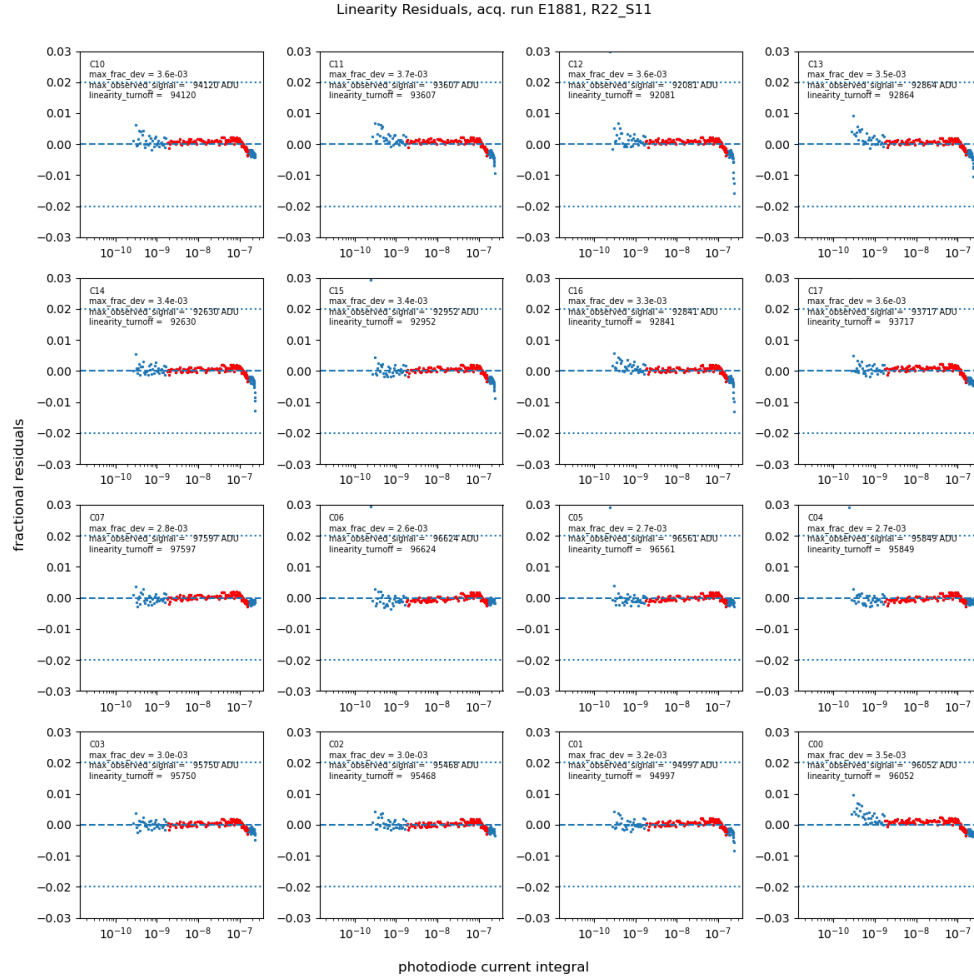


Figure 154: Linearity residual plots for R22_S11 from run E1881.

A.2.3.4 Data from flat pairs For studies of row-means variance, a plot showing the row-to-row mean vs. variance of differences between a pair of flats is created. This plot shows different amplifiers in different colors, and allows for identification of outliers of row-wise correlated noise for individual amplifiers (Fig. 155).




Figure 155: Row means variance plot for R22_S11 from run E1880.

A.2.3.5 Charge-transfer inefficiency Plots showing the dependence of charge-transfer inefficiency (CTI) on signal are also generated. Different amplifiers are plotted in different colors. Plots are generated for both serial and parallel CTI; Figure 156) shows an example serial CTI plot.




Figure 156: Serial CTI plot for R22_S11 from run E1880.

A.2.3.6 BF covariances From the flat pair acquisitions, plots measuring the brighter-fatter covariance as a function of flux are generated. Amplifiers are plotted in different colors so that outliers can easily be traced to their respective sensors (Fig. 157).

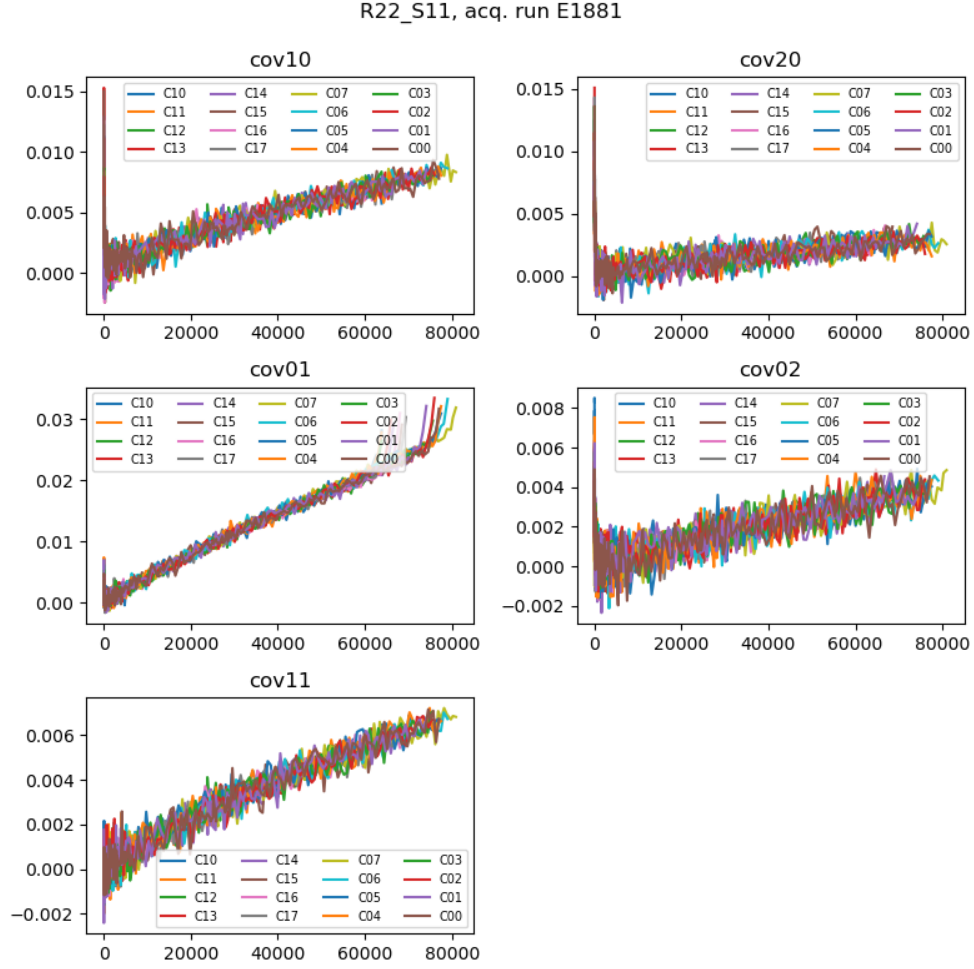


Figure 157: Brighter-fatter covariance plot for R22_S11 from run E1880.

A.2.3.7 Persistence B protocol runs include a persistence acquisition, which is a saturated flat (usually at 400k e-/pixel), followed by a sequence of multiple 15 s dark images. From analysis of this image acquisition sequence, persistence plots are generated that show the decay of residual signal in the dark sequence. These plots are color coded according to amplifier (Fig. 158).




Figure 158: Persistence plot for R22_S11 from run E1880. Units are ADU.

B Observatory Control System integration

A secondary goal of the Run 7 campaign was testing integration between the Camera Control System (CCS) and the Observatory Control System (OCS). The tests performed with OCS control of LSSTCam during this campaign are the first operations of the LSSTCam with OCS, which will be the normal operating mode during operations of the Simonyi Survey Telescope. To this end, many tests were performed during the Run 7 campaign including basic functionality tests, short OpSim runs, long, scheduler-driven OpSim runs (“survey mode”) and the start of mock calibrations.

B.1 LSSTCam in OCS

B.1.1 Configurable SAL Components: base unit of the OCS

Every component controlled by the observatory is represented by a Configurable SAL Component (CSC) in the control system. CSCs are remote objects capable of receiving command and publishing telemetry and events. During the Run 7 campaign, CSC communicated with DDS middleware, which was replaced by Kafka middleware in early 2025. Most CSCs interact with lower-level software such as an engineering interface, pipelines, or in the case of all cameras (LSSTCam, ComCam and LATISS) the ocs-bridge component of CCS.

LSSTCam is represented by a group of three CSCs:

1. MTCamera: this is the ocs-bridge in CCS as seen by the OCS – not exactly a CSC in the traditional sense, but often called a CSC. ocs-bridge communicates over DDS/Kafka middleware just as any CSC would.
2. MTOODS: Main Telescope Observatory Operations Data Service – handles data ingestion to LSST butler and purging old data from disks once ingested.
3. MTHederService: manages FITS header data for images (science program, scheduler config, weather data and more)

These can all be seen on LOVE, LSST Observatory Visualization Environment, ASummaryState view under the LSSTCam group of CSCs in Figure 159.




Figure 159: LSSTCam CSCs in LOVE.

CSCs in general have four regular operation states and two other states with the general rules. Normal states:

1. STANDBY (transitions only to DISABLED): software is running on OCS computing clusters, but is not doing anything.
2. DISABLED (Transitions to STANDBY, ENABLED, FAULT): Software is running and communicating with hardware. Telemetry is sent to the EFD.
3. ENABLED (Transitions to DISABLED, FAULT): Software is running and ready to receive commands.
4. FAULT (transitions only to STANDBY): Software encountered a critical issue. Watcher alarms will be activated.

Extra states:

1. OFFLINE: The software is not available and active within the OCS computing clusters.
2. UNKNOWN: The software is not deployed within the OCS computing clusters.

However, MTCamera, ATCamera and CCCamera are not exactly CSCs, but rather the OCS-bridge. OCS-bridge maintains all states, but practically speaking has different behavior for only three of them:

1. OFFLINE: OCS-bridge is not being used by the OCS. This is the state it should be in when controlling a camera through lower-level CCS commands.
2. STANDBY: OCS-bridge has control of camera systems. The OCS itself cannot issue commands to the camera in this state, however. This is a “safe mode” during the day and end of the night.

3. **ENABLED:** OCS-bridge has control of camera systems. The OCS can issue commands to a camera. The camera can be controlled through scripts and the scheduler. This is the state when the camera is in use for calibrations and operations.

Telemetry is always sent to the EFD from CCS.

B.1.2 ts_xml

Both the software interface definitions and official ICD of T&S/OCS systems are handled by `ts_xml`. The human-readable ICD version of `ts_xml` may be found here: . The GitHub management of these files (+ machine readable versions) is found here:

.

As the camera and ocs-bridge interface with OCS, the CCS team generates an XML file that defines all commands, events and telemetry that the OCS knows about from the ocs-bridge for each CCS release. This can be found under MTCamera (also ATCamera and CCCamera for the other cameras) here (human-readable): and here (machine-readable):

B.1.3 OCS Scheduler and Script Queue

OCS is operated via python scripts, which can be seen in these repositories: Standard Scripts and External Scripts. Standard scripts undergo rigorous testing to ensure functionality across all versions. These form the basic functions of the observatory. External scripts have less rigorous testing and are not guaranteed to work across versions. These are for specific, sometimes one-off tests of observatory systems.

Scripts to be executed are added to a script queue for the observatory. At the time of writing, three script queues are available: AuxTel, MainTel and OCS. AuxTel and MainTel are for operating their respective telescopes, the OCS queue is a backup queue for other operations, for example, LSSTCam EO testing.

Aside from manual addition of scripts to the queue, each script queue has an associated scheduler which may add scripts to the queue through predefined scripts (JSON blocks) or programmatically with the Rubin Scheduler taking input data such as time of day, weather, fea-

sibility and more based on a scheduler configuration. Configurations and JSON blocks may be found under the Scheduler section of `ts_config_ocs`.

Both the script queues and their respective schedulers are represented through CSCs. SAL index 1 corresponds to MainTel, 2 corresponds to AuxTel and 3 corresponds to OCS.

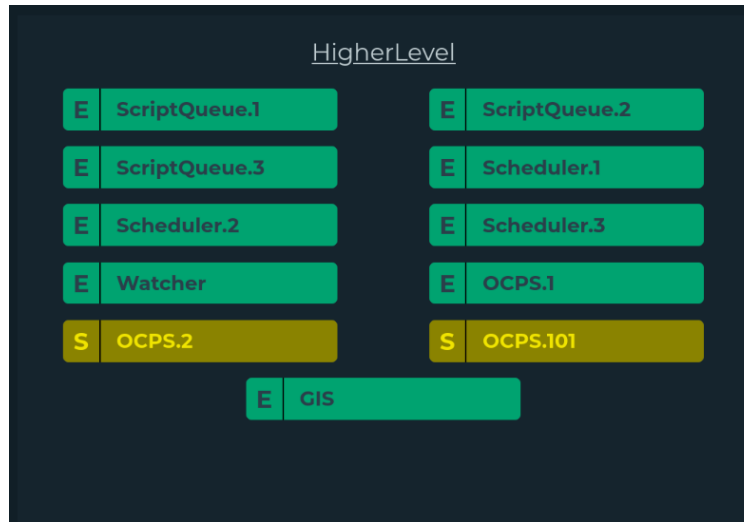


Figure 160: Visualization of scripts= queues and schedulers in LOVE.

B.1.4 OCS-Bridge

ocs-bridge is a CCS subcomponent that communicates with the OCS. Instead of being a standard CSC, it has its own interface to the observatory via DDS (soon to be kaftka) middleware. It issues commands to the camera via MCM. It also raises events upon command/activity completion. These events are used by the OCS/scripts to determine when to stop waiting for an action to complete.

An example procedure for taking control of the camera with OCS /ocs-bridge can be viewed in this test case: BLOCK-T154

ocs-bridge must be in offline-available to allow the OCS to take control. This is done with ocs-bridge setAvailable with a lock on ocs-bridge in a ccs-shell. For OCS to successfully bring MTCamera to enabled (ocs-bridge in ENABLED state) MCM must be able to transition its controlled subsystems to normal mode. This requires a few conditions, some of which are listed here:

- MCM must be able to enter normal mode
- Shutter must be closed
- Alarms on subsystems such as focal-plane must be cleared

B.2 Testing steps and achievements in Run 7

OCS testing with LSSTCam was part of three steps of the Run 7 campaign. Prior to starting any of the steps, basic functionality was confirmed with single exposures via an OCS script. The first image taken with OCS was on September 13, 2024, early in the campaign to give time to address potential issues. The three steps with OCS-driven operation were then as follows:

- 30 min OpSim runs. These were pre-generated sequences pulled from survey simulations including filter changes and different observation strategies for different filters, which may be a feature of the LSST survey. Two runs were successfully completed: E1414 which included operation of the shutter and E1585 which did not operate the shutter.
- Fully scheduler driven OpSim runs. These runs used the Rubin Scheduler to dynamically choose targets, filter changes and exposure times depending on the filter. The scheduler was configured to choose targets at any time - not only with the Sun down - and to only operate LSSTCam and not point the telescope, though simulated slew times were added.
- A run of an example calibration sequence that could be part of a daily calibration sequence. This was intended to include testing image ingestion and processing of the calibrations with the OCS.

Together, these steps demonstrated the ability to operate LSSTCam in a mode similar to what nighttime operations are anticipated to be like in a steady state. The listed acquisitions include more than 12 hr of operation under control of the OCS including filter changes, shutter operation, guider ROI setting and variable exposure times. These operations also included simulated slew times. Image acquisition also worked in conjunction with the observatory header service to populate image headers with additional details and telemetry from the observatory.

While testing was largely successful, some issues were encountered and resolved along the way. One issue was that the interface for sending guider regions of interest (ROIs) through the OCS was incomplete. The initial interface was developed and tested as part of this process.

The sending of guider ROIs revealed another issue with command acknowledgments between the camera OCS-bridge and the scripts operating the system.

B.3 Further testing and development

Not all testing goals to verify the Camera-OCS interface were possible during Run 7. A future testing period for lingering items is planned in February or March, 2025 to close the loop on the remaining tests. These tests will use playlists and playback of data taken during Run 7.

An item to test is the summit storage facility for operations (S3), which was not yet available during Run 7. With the playback mode testing, we will verify proper ingestion of images into S3.

Another open item to be tested is a complete loop of calibrations being processed with the Observatory Control Pipeline Service (OCPS) immediately after their acquisition. A calibration playlist is included as part of the planned OCS testing period. With this, data can be collected and sent to the OCPS to verify near real time processing of calibrations.

Lastly, between Run 7 and the first photon of LSSTCam, the observatory changed the communication middleware software used within the OCS. The change from DDS to Kafka is expected to improve stability of the communication between distributed software components of the OCS. For LSSTCam, the playback testing will be used to test both general performance in the new middleware environment, and whether missing command acknowledgments found in Run 7 remain reproducible.

C Phosphorescence

Because we have discovered the phosphorescence in ITL sensors only recently, we have here cobbled together the data products described here. We believe they are appropriate for quantitative characterization of the phosphorescence systematics that are essentially a position-dependent hysteresis indexed by recent illumination history.

Happily, the *front-side persistence* seen in e2v sensors has been strongly suppressed with only parallel clock swing voltage adjustment (see Fig. 158 for post-mitigation performance). Although the working assumption regarding the *e2v persistence* was that it affected all pixels identically, this has not been demonstrated, and only sensor-averaged *persistence* expression has been evaluated to date. The current, post-mitigation *persistence* performance (Fig 84) provides indirect evidence of the spatial dependence of the *persistence*, revealed in the *skew* of the distributions quantified:⁸ Especially with *y-band* illumination, the significant skew apparently indicates pixel-to-pixel variation in persistence, which has not yet received adequate attention as a mechanism that would fall within ISR scope.

The *phosphorescence* seen in ITL sensors is highly position dependent as compared to the *front-side persistence*. This was first revealed plainly in the morphology of post-illumination dark images of R43_S11 (see Fig. 137f) and R00_SW1 (see Fig. 135). Clearly, the methodology of *persistence* estimation feedback, used while tuning clock swing, is not suitable for quantifying *phosphorescence* in parts of the focal plane covered by ITL sensors.

The data products displayed in the following sections are described here:

- * Identification of phosphorescent regions (Appendix D): Using the 20 B-protocol runs containing the *BOT PERSISTENCE* set of images, we constructed pixel-by-pixel median images over the 20 instances of the *first* post-flash **dark** image, and also for the *twentieth* post-flash **dark** image accumulated for each of the runs. The difference between the two (*first* minus *twentieth*) is 8x8 blocked and displayed on a raft-by-raft basis. Such pixels appear significant when they contain (on average) $10e^-$ of transient signal within the first 15 s dark.
- * Comparison against *blue* flat field response (Appendix E): The same *transient term* images

⁸NB. Disentangling true *persistence* from residual charge read out following *clear* operations with finite inefficiency remains difficult. However, it is believed that operating with **NopSf** removes nearly all residual signal resulting from incomplete clear.

described above are displayed *without* blocking, no rebinning. These were inspected alongside the *blue* flat field response images to estimate whether *coffee stain* and *vampire pixel* features are consistently correlated with phosphorescent regions. (No for *coffee stains*, yes for *vampire pixels*.)

- * Kinetics characterization (Appendix F): Pixel-by-pixel median images constructed with the 20 *BOT PERSISTENCE* datasets were constructed for each of the 20 successive dark images following the flash (trigger) exposure. These data products did not have the twentieth median image subtracted from each, but were kept intact. Using ROI pixel extractions and signal level quantile estimations, we extract high-significance evolution of the transient signal, with the understanding that different pixels may be driving each time slice's quantile level. In principle, this could be repeated for each pixel in the focal plane to generate a short list of parameters modeled after the function given in Fig. 179⁹. Evolution of the actual cumulative distributions for a selection of ROIs is provided.
- * Response characterization (Appendix G): The acquired data available to address phosphorescence response to illumination (wavelength and signal level in the trigger as independent variables) were marginally adequate to estimate the *amplitude only* of a putative decay profile resembling Figure 179. We did not have 20 *BOT PERSISTENCE* datasets from which to construct median images, and so worked only with the single instances available (specifically, only the first dark image following the flash). Variations from ROI to ROI in the phosphorescence response are clearly seen across the independent variable plane (illumination wavelength and signal level)¹⁰.

⁹These results are expected to be sensitive to details of the *trigger* exposure and where it is placed temporally within the exposure duration. We currently do not have such transient functions measured after illumination exposure defined by the LSSTCam Shutter's *open* duration.

¹⁰Note that ROI extractions from single dark images such as these are susceptible to cosmic rays. The 99% quantile level applied to the extraction also gets pegged to the noise distribution at about 3σ when the product of phosphorescent response and trigger illumination level is small.

D Phosphorescence identification on a set of ITL sensors




Figure 161: Phosphorescence transients for the R00 CRTM captured in the first 15 s following *red* CCOB LED at $400 \text{ ke}^-/\text{pixel}$. With 8×8 blocking, the upper end of the color scale (640) corresponds to $10 \text{ e}^-/\text{pixel}$ when averaged over 64 pixels contributing.




Figure 162: Phosphorescence transients for the R01 RTM captured in the first 15 s following *red* CCOB LED at $400 \text{ ke}^-/\text{pixel}$. With 8×8 blocking, the upper end of the color scale (640) corresponds to $10 \text{ e}^-/\text{pixel}$ when averaged over 64 pixels contributing.




Figure 163: Phosphorescence transients for the R02 RTM captured in the first 15 s following *red* CCOB LED at $400 \text{ ke}^-/\text{pixel}$. With 8×8 blocking, the upper end of the color scale (640) corresponds to $10 \text{ e}^-/\text{pixel}$ when averaged over 64 pixels contributing.




Figure 164: Phosphorescence transients for the R03 RTM captured in the first 15 s following *red* CCOB LED at $400 \text{ ke}^-/\text{pixel}$. With 8×8 blocking, the upper end of the color scale (640) corresponds to $10 \text{ e}^-/\text{pixel}$ when averaged over 64 pixels contributing.




Figure 165: Phosphorescence transients for the R04 CRTM captured in the first 15 s following *red* CCOB LED at $400 \text{ ke}^-/\text{pixel}$. With 8×8 blocking, the upper end of the color scale (640) corresponds to $10 \text{ e}^-/\text{pixel}$ when averaged over 64 pixels contributing.




Figure 166: Phosphorescence transients for the R10 RTM captured in the first 15 s following *red* CCOB LED at 400 ke^- . With 8×8 blocking, the upper end of the color scale (640) corresponds to $10 \text{ e}^-/\text{pixel}$ when averaged over 64 pixels contributing.




Figure 167: Phosphorescence transients for the R20 RTM captured in the first 15 s following *red* CCOB LED at $400 \text{ ke}^-/\text{pixel}$. With 8×8 blocking, the upper end of the color scale (640) corresponds to $10 \text{ e}^-/\text{pixel}$ when averaged over 64 pixels contributing.




Figure 168: Phosphorescence transients for the R40 CRTM captured in the first 15 s following *red* CCOB LED at $400 \text{ ke}^-/\text{pixel}$. With 8×8 blocking, the upper end of the color scale (640) corresponds to $10 \text{ e}^-/\text{pixel}$ when averaged over 64 pixels contributing.




Figure 169: Phosphorescence transients for the R41 RTM captured in the first 15 s following *red* CCOB LED at $400 \text{ ke}^-/\text{pixel}$. With 8×8 blocking, the upper end of the color scale (640) corresponds to $10 \text{ e}^-/\text{pixel}$ when averaged over 64 pixels contributing.




Figure 170: Phosphorescence transients for the R42 RTM captured in the first 15 s following *red* CCOB LED at $400 \text{ ke}^-/\text{pixel}$. With 8×8 blocking, the upper end of the color scale (640) corresponds to $10 \text{ e}^-/\text{pixel}$ when averaged over 64 pixels contributing.




Figure 171: Phosphorescence transients for the R43 RTM captured in the first 15 s following *red* CCOB LED at $400 \text{ ke}^-/\text{pixel}$. With 8×8 blocking, the upper end of the color scale (640) corresponds to $10 \text{ e}^-/\text{pixel}$ when averaged over 64 pixels contributing.




Figure 172: Phosphorescence transients for the R44 CRTM captured in the first 15 s following *red* CCOB LED at $400 \text{ ke}^-/\text{pixel}$. With 8×8 blocking, the upper end of the color scale (640) corresponds to $10 \text{ e}^-/\text{pixel}$ when averaged over 64 pixels contributing.

E Phosphorescence morphological comparisons with features seen in *blue* flat field response

Figures 173 through 178 are an incomplete selection of ITL sensors with phosphorescence. They compare expressed phosphorescence (transient term) with the *blue* CCOB LED flat response. Inspection of these images would lead one to conclude that in certain cases, the phosphorescence patterns resemble the coffee stain patterns' regions of lower QE at short wavelength (*cf.* Fig. 135, Fig. 177). In other cases, the opposite appears to be true (*cf.* Fig. 174, Fig. 175). In several cases, there appear to be no particular correlations.

In cases where variations in the blue flat-field response are due to *vampire pixels* (*cf.* Fig. 136) with a completely different wavelength dependence, presumably due to depth dependence in the direction of the drift field lines), we see *high amplitude* and *long timescale* transient phosphorescence associated with these vampire pixel complexes. These tend to be the brightest phosphorescent features we see, and this fact may provide a strong hint regarding the origin of this phosphorescence phenomena. These quantitative differences are most easily seen in the kinetics discussion, Section F.




Figure 173: The ITL sensor R01_S00. Top: the transient phosphorescence term. Bottom: the *blue* flat response. The large, extended spot appears to be centered on a vampire pixel, which also expresses a large amplitude of phosphorescence, which emits enough current to contaminate the parallel overscan in at least the first 15 s exposure following trigger. The flat response feature has opposite polarity from the phosphorescence.




Figure 174: The ITL sensor R02_S02. Top: the transient phosphorescence term. Bottom: the *blue* flat response. The *coffee stain* feature in the flat response has the same polarity as the phosphorescence. A phosphorescent vampire pixel is seen in segment R02_S02_C07.

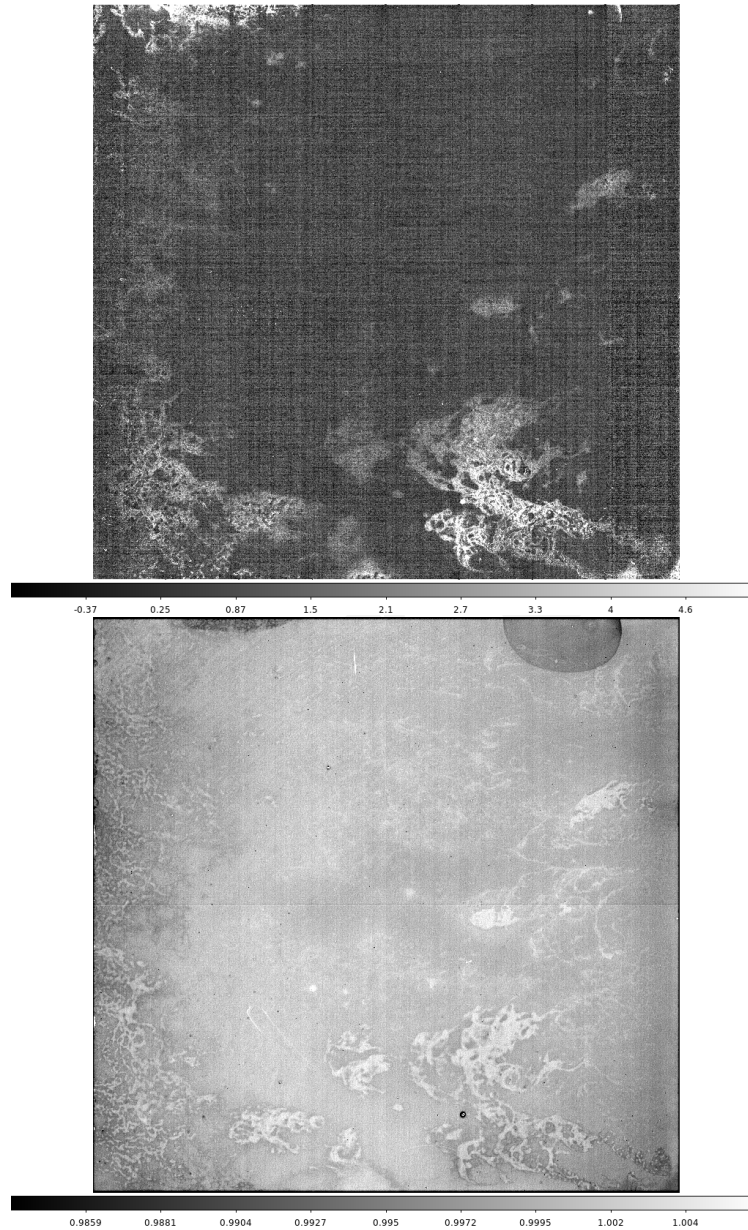


Figure 175: The ITL sensor R02_S12. Top: the transient phosphorescence term. Bottom: the *blue* flat response. Generally the polarity of the phosphorescence matches that of the *coffee stain* in the flat field response, but exceptions include the vampire pixel seen in segment R02_S12_C05.




Figure 176: The ITL sensor R03_S10, detail of the vampire of R03_S10_C15. Top: the transient phosphorescence term. Bottom: the *blue* flat response. As in previous examples, this vampire pixel's transient term is large enough to contaminate the parallel overscan even after the first 15 s following trigger.




Figure 177: The ITL sensor R43_S11. Top: the transient phosphorescence term. Bottom: the *blue* flat response. This sensor appears to have the largest integrated phosphorescence among ITL sensors studied. The flat response feature has opposite polarity from the phosphorescence.




Figure 178: The ITL sensor R43_S20, segments C00 through C03. Top: the transient phosphorescence term. Bottom: the *blue* flat response. This sensor apparently exhibits peculiar radial crazing patterns seen in both phosphorescence as well as in flat field response, with polarities aligned.

F Phosphorescence kinetics characterization

Figures 180 through 186 quantify the expressed phosphorescence distributions in ROIs on seven of the problematic ITL sensors. Previously, we had captured the phosphorescence *transient term* across the ITL sensors (cf. Figs. 161 thru 172); here we track ROI pixel distribution parameters of individual median images constructed from the selection of specific images acquired across the 20 B-protocol datasets available (listed in Table 13).

By fitting decay models to these persistence curves, it is immediately clear that there are multiple (>2) timescales at play for the pixels in each ROI. An example of such a fit is given in Figure 179 where a 3-population relaxation model is used to characterize evolution of the 99% quantile level of the distribution. In this case, there are three different exponential timescales determined: $(\tau_1, \tau_2, \tau_3) = (0.62, 2.5, 18.3)$ in image units (10.9, 43.8 and 320 s, respectively). The corresponding ratio of these populations works out to 4.5% (fast), 21.5% (medium) and 74% (slow), respectively. Inspection of the more detailed parameters plotted generally indicate skewed distributions from mismatches between medians and means; the choice of the 99% quantile level to characterize was mainly to estimate the degree to which images would need to be phosphorescence-corrected (and/or the variance plane modified, given the asymmetric impact of the position-specific phosphorescence contribution in recorded images).




Figure 179: A three-population fit of the phosphorescence expressed by the vampire pixel region of R20_S20_C13. The fit was performed on the 99% quantile level where signal levels are well above the 3σ level of the noise distribution. Here, image numbers are parasitically used as time units, with roughly 17.5 s per image.

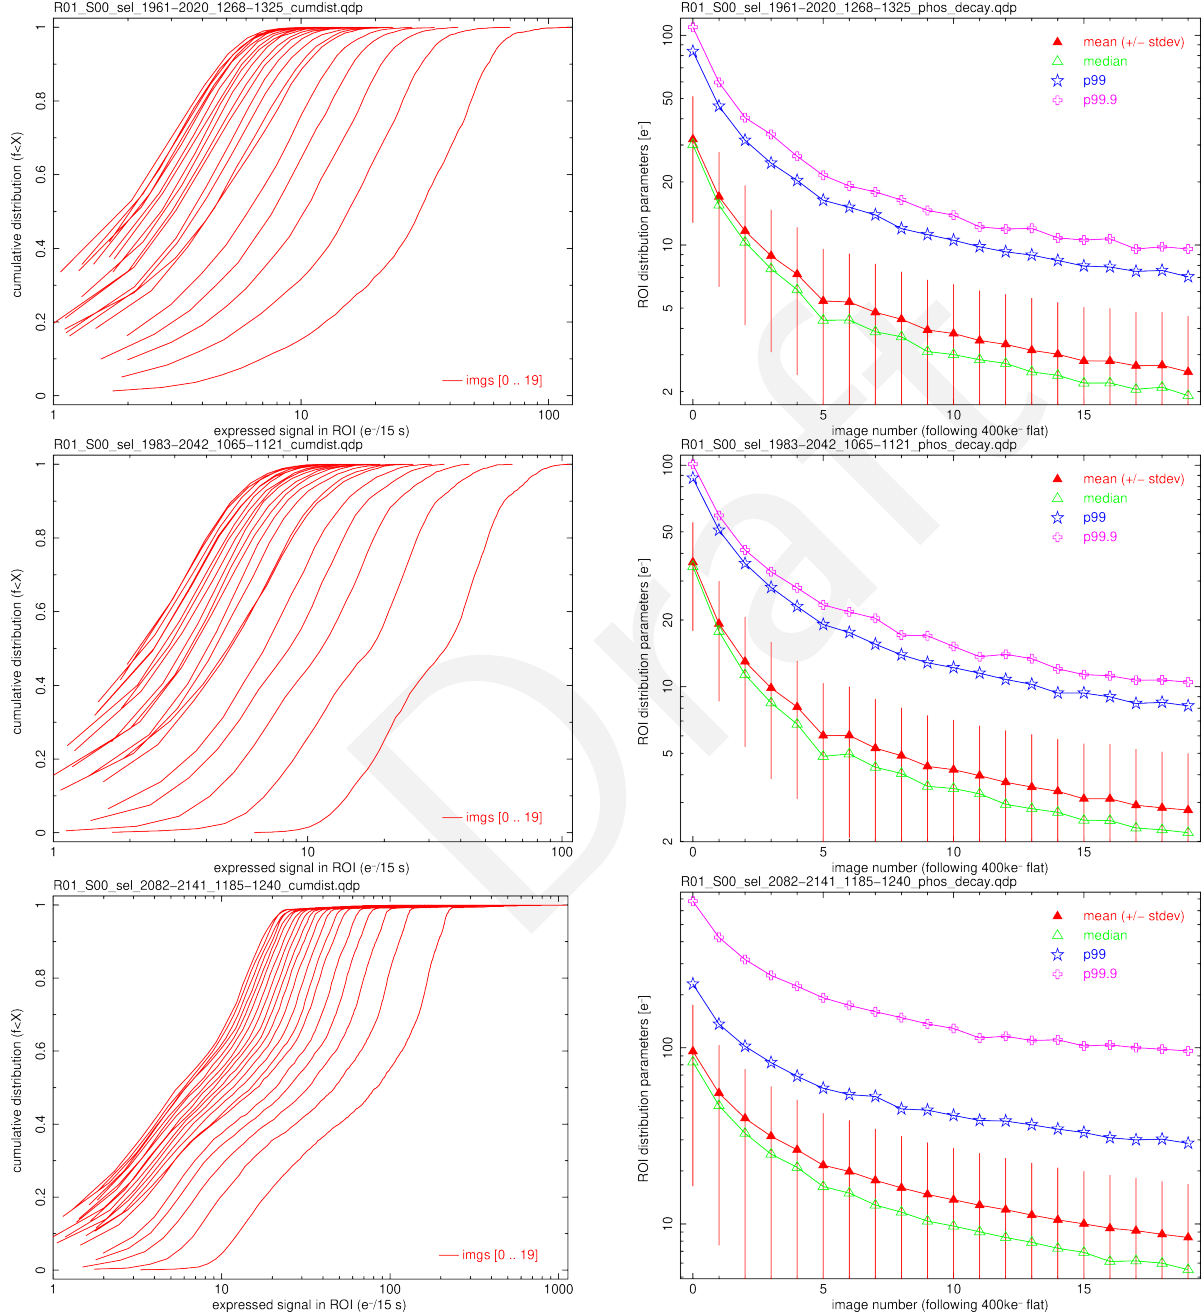


Figure 180: Kinetics for phosphorescence expression in ROIs of images for R01_S00. This is the prominent cosmetic seen in Fig. 173, which is apparently a vampire pixel.

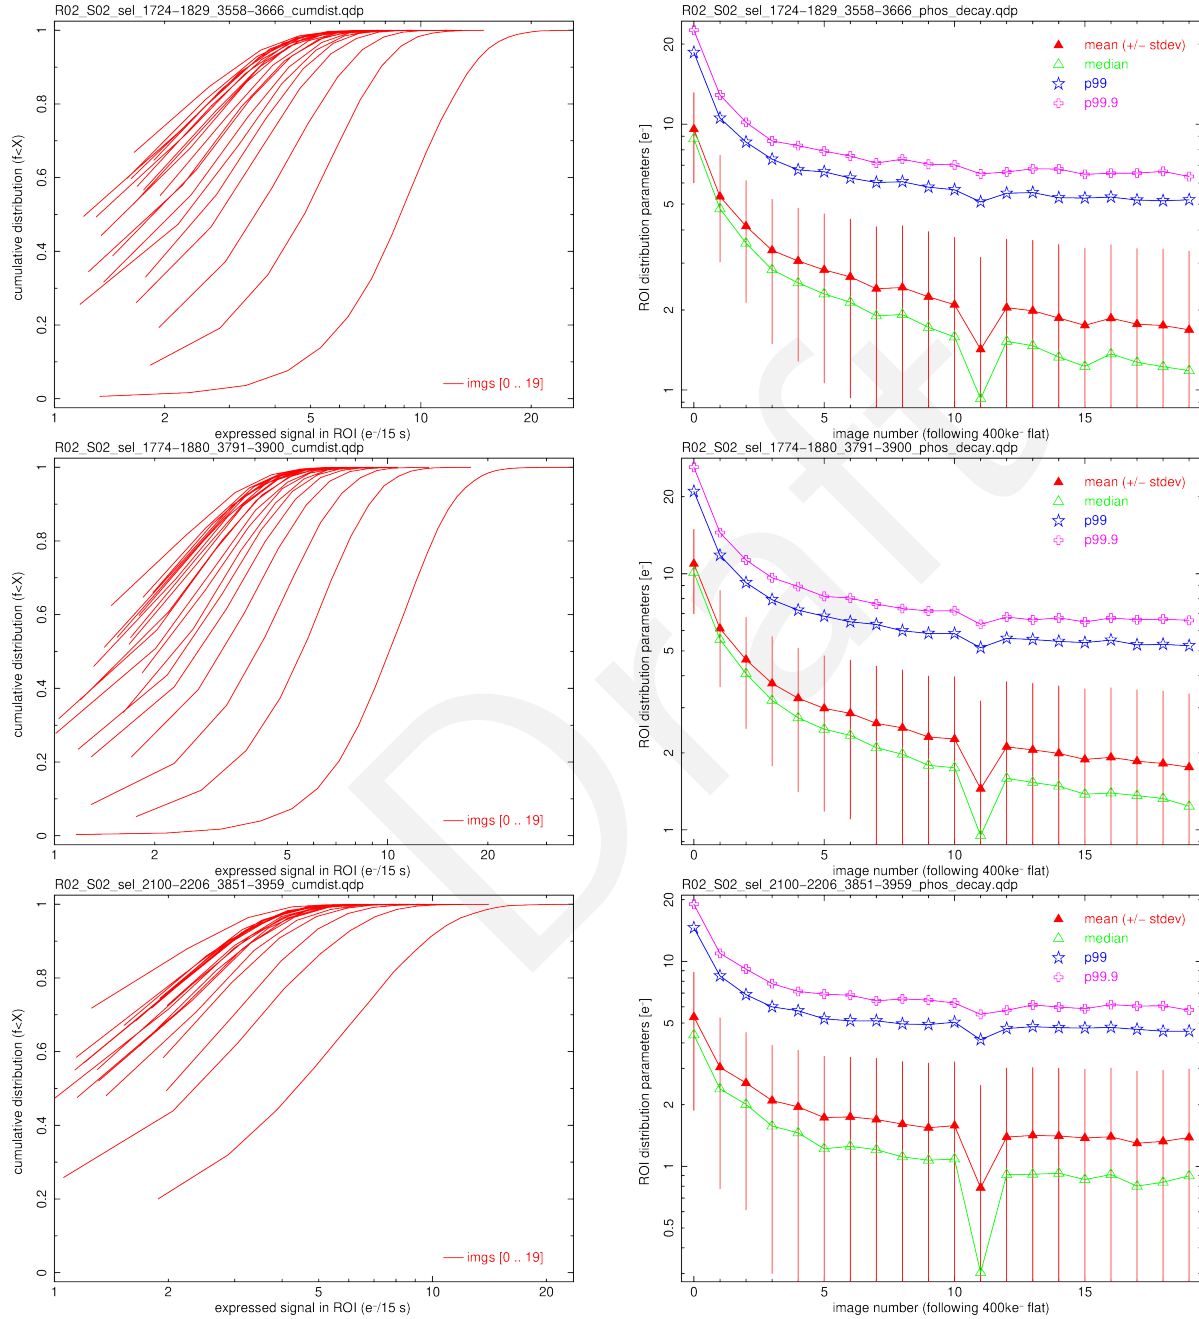


Figure 181: Kinetics for phosphorescence expression in ROIs of images for R02_S02. This is the diffuse phosphorescence that correlates with the coffee stains seen in Fig. 174. No extractions were performed on the vampire pixels found on the same sensor (R02_S02_C15 and R02_S02_C07).

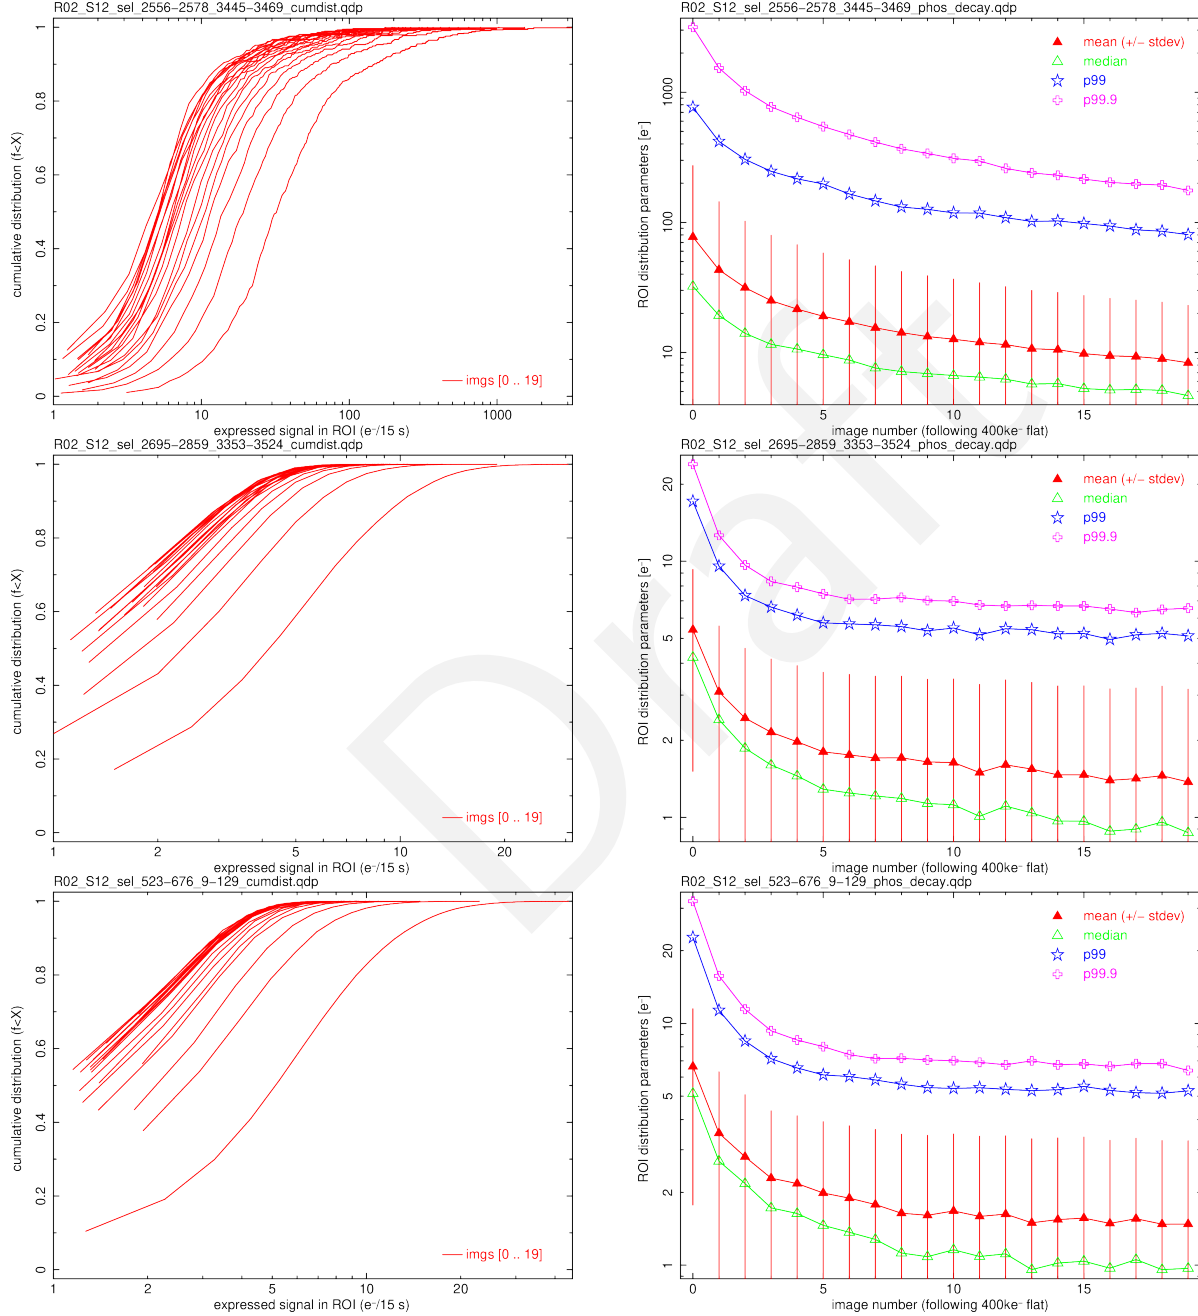


Figure 182: Kinetics for phosphorescence expression in ROIs of images for R02_S12. This is the structured phosphorescence that correlates with the coffee stains seen in Fig. 175.

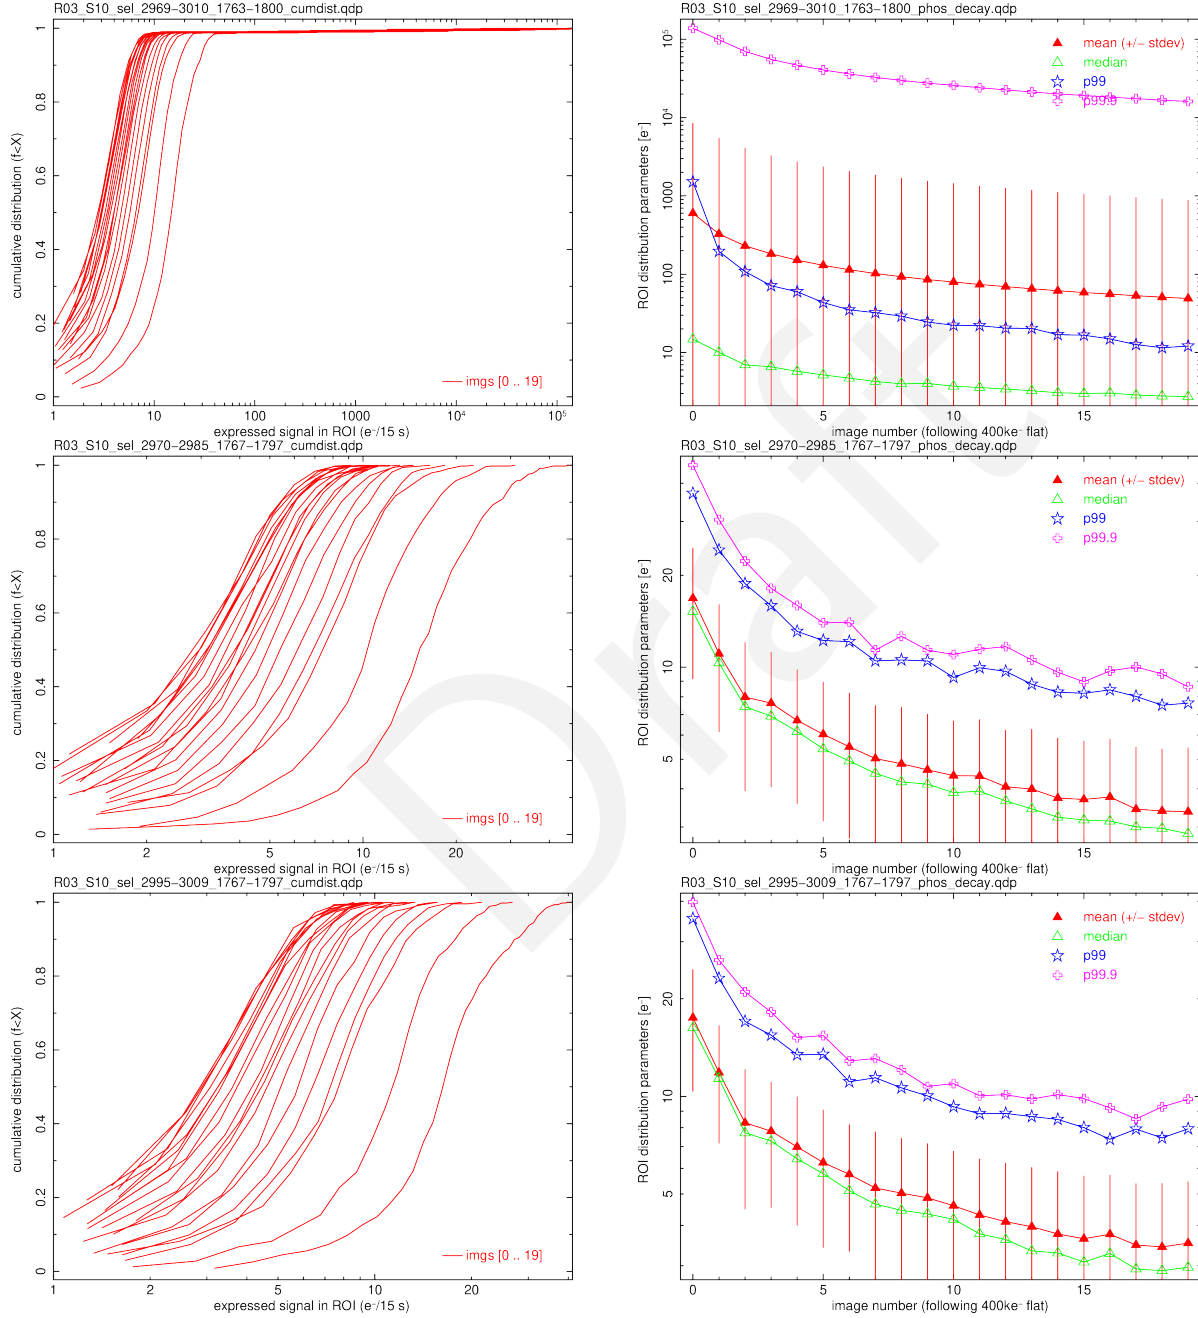


Figure 183: Kinetics for phosphorescence expression in ROIs of images for R03_S10. These describe regions including or near the bright/focusing vampire pixel seen in Figs. 176, 136b and 137b.




Figure 184: Kinetics for phosphorescence expression in ROIs of images for R20_S20. These describe the prominent non-focusing vampire pixel seen in Figs. 136d and 137d.

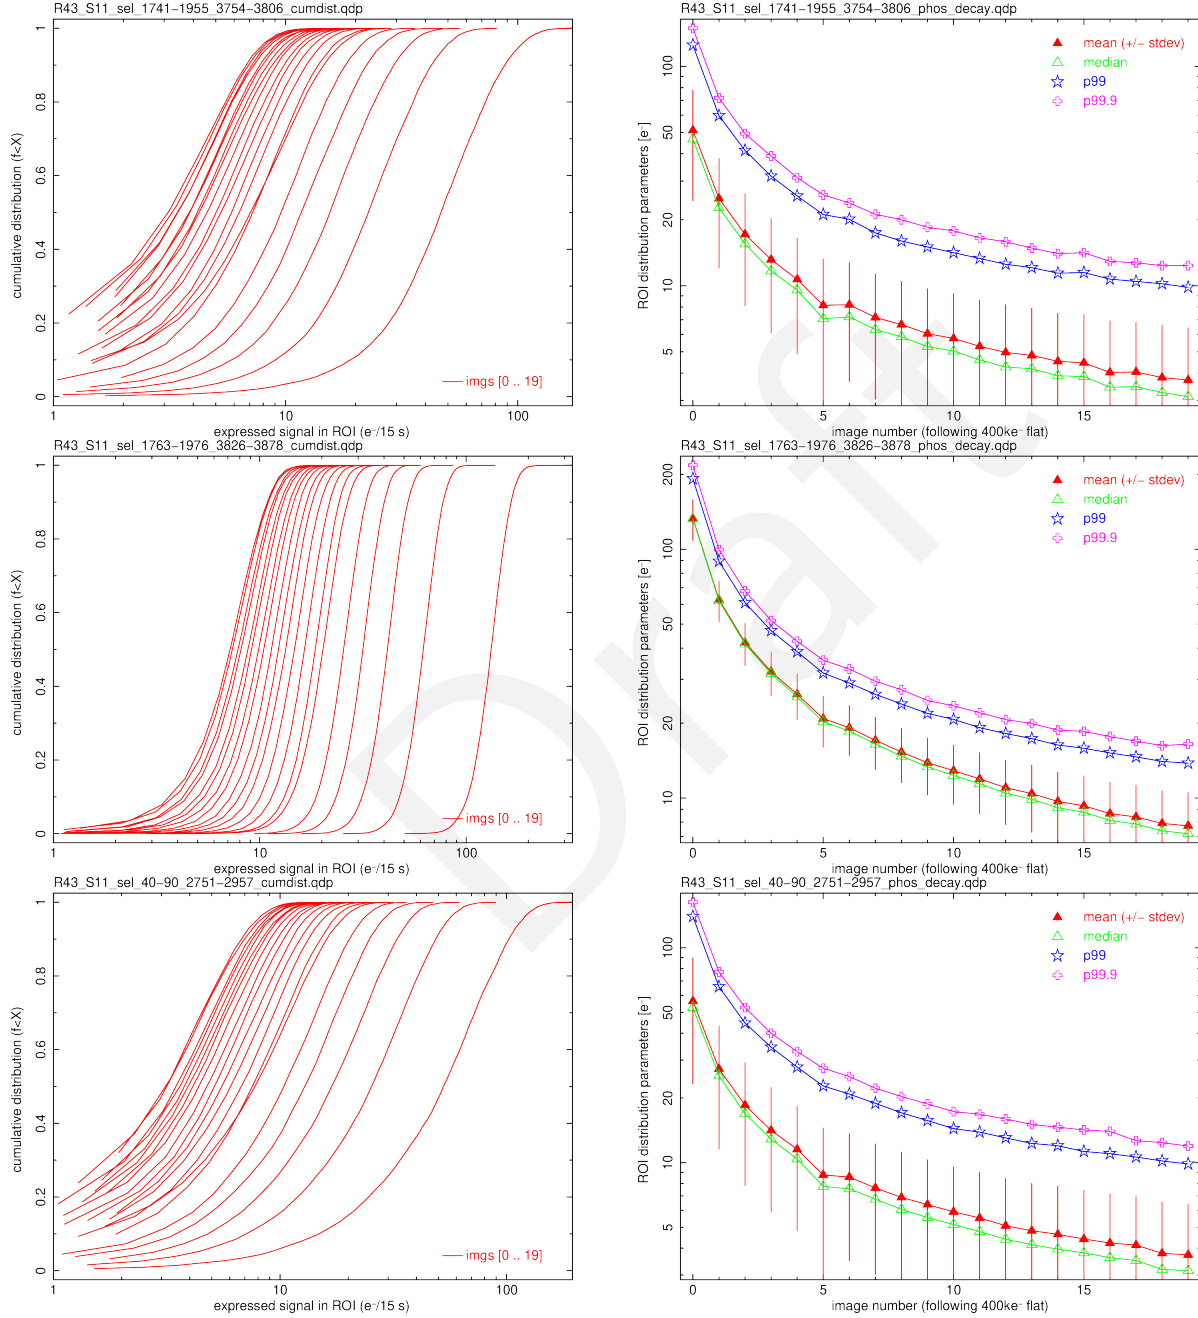


Figure 185: Kinetics for phosphorescence expression in ROIs of images for R43_S11. These describe bright, diffuse transient regions seen in Figs. 177 and 137f, which apparently turn off completely when the HV Bias is *off*.

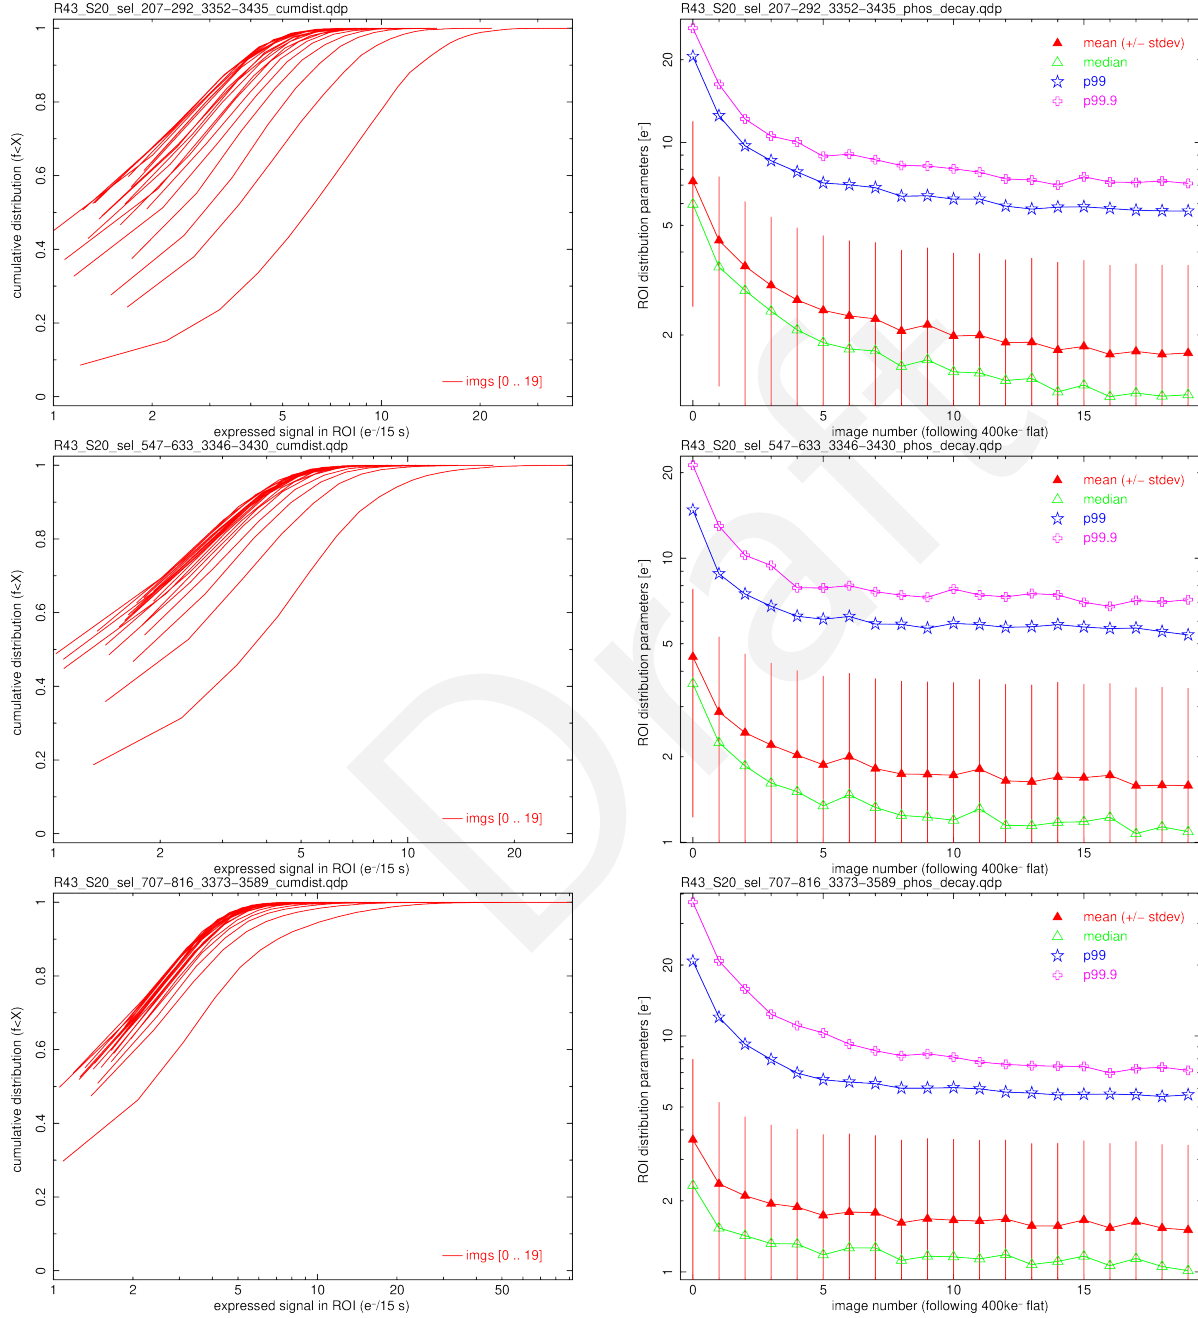


Figure 186: Kinetics for phosphorescence expression in ROIs of images for R43_S20. These include some of the the highly structured, snowflake-like transient regions seen in Figs. 137h and 178.

G Phosphorescence response characterization

Figures 187 through 193 attempt to quantify the expressed phosphorescence response in ROIs on seven of the problematic ITL sensors. Previously, we had captured the phosphorescence *transient term* across the ITL sensors (*cf.* Figs. 161 through 172); we also tracked ROI pixel distribution parameters of individual median images constructed from the selection of specific images acquired across the 20 B-protocol datasets available (listed in Table 13). Here we analyze the signal level- and wavelength-dependencies of the expressed phosphorescence captured in the first dark image following a flat exposure. Table 28 provides the dataset IDs and SeqIDs used for this analysis.

Table 28: Zephyr Scale E-numbers and corresponding SeqIDs analyzed to estimate signal level and wavelength dependence of the phosphorescence response.

Run numbers and SeqIDs of first dark following trigger		
CCOB LED	trigger flat target signal	runID & SeqID
uv	500	E1770:20241028_000010
uv	1000	E1503:20241020_000489
uv	1500	E1771:20241028_000039
uv	3000	E1504:20241020_000533
uv	4500	E1772:20241028_000068
uv	5000	E1505:20241020_000577
uv	10000	E1506:20241020_000621
uv	13500	E1773:20241028_000097
uv	30000	E1507:20241020_000665
blue	500	E1774:20241028_000126
blue	1000	E1502:20241020_000445
blue	1500	E1775:20241028_000155
blue	3000	E1501:20241020_000401
blue	4500	E1776:20241028_000184
blue	5000	E1500:20241020_000357
blue	10000	E1499:20241020_000313
blue	13500	E1777:20241028_000213
Continued on next page		

Table 28 – continued from previous page

Run numbers and SeqIDs of first dark following trigger		
blue	30000	E1498:20241020_000269
blue	50000	E1491:20241018_000989
blue	150000	E1485:20241018_000725
blue	400000	E1484:20241018_000678
red	50000	E1490:20241018_000945
red	150000	E1486:20241018_000769
red	400000	E1483:20241018_000634
nm750	50000	E1492:20241018_001033
nm750	150000	E1487:20241018_000813
nm750	400000	E1479:20241018_000543
nm850	50000	E1493:20241018_001077
nm850	150000	E1488:20241018_000857
nm850	400000	E1477:20241018_000455
nm960	50000	E1494:20241018_001121
nm960	150000	E1489:20241018_000901
nm960	400000	E1478:20241018_000499

These runs were performed to sample a two-dimensional, rectangular parameter space, and each measurement was executed only once. The resulting sampling was completed incrementally, over 3 separate days. Using only one image for each data point, we were not able to median multiple images acquired under identical conditions (as we had done for Sections D, E and F).

Because there is significant variation in morphological characteristics of the phosphorescence, we adopted the following strategy to quantify phosphorescence expression in each: Once the image is processed through ISR, each of the sensor-specific ROIs is used to filter the pixels, and the signal distribution parameters are evaluated. The 99% quantile signal level was used as a bright-end proxy for the expressed phosphorescence for each ROI. A consequence of this

is that when there is insignificant or undetectable phosphorescence, this proxy choice would be artificially high, which would be pegged at about 4σ above the noise distribution mean.

The figures provide dashed lines that represent constant *phosphorescence efficiency ratios* to guide the eye (at 10%, 1% and 0.1%), while different color LED illuminations are represented by different plotting symbols and line colors. The only CCOB LED used across the entire range of trigger flat signal levels is the *blue* one. The $400ke^-/\text{pixel}$ *blue* LED trigger flat-induced phosphorescence levels are the only ones described thus far in Sections D, E and F.

Two of the sensors exhibiting distributed and structured phosphorescence expression (R02_S02 & R43_S20) appear to have phosphorescence yields below 3×10^{-3} for *uv* and *blue* LED illumination. Given the kinetics studied for these ROIs for *blue* LED illumination (*cf.* Figs 181 and 186), the worst-case contribution may be $55 e^-/\text{pix}/15\text{s}$ (*uv* LED, $30ke^- \times e^{+0.62} \times 10^{-3}$). The $e^{+0.62}$ scaling factor comes from the fact that the LED flash occurs (and ends) at the beginning of the trigger flat illumination and typically lasts for just a fraction of the image time of $\sim 17.5\text{ s}$.

One other sensor with distributed, coffee stain-like phosphorescence (R02_S12) shows significantly more signal in one of the ROIs for (*uv* and *blue* LEDs for $30ke^-/\text{pixel}$ and $400ke^-/\text{pixel}$, respectively). Upon applying the $e^{+0.62}$ factor, the worst case phosphorescence here would scale to $560e^-/\text{pixel}/15\text{s}$ and $1800e^-/\text{pixel}/15\text{s}$, respectively.

The remaining four sensors (R01_S00, R03_S10, R20_S20 & R43_S11) show even more phosphorescence. The first three of these are due to vampire pixels (with or without central hot spots), while the last one showed diffuse glow along the edges that “shuts off” with HV Bias. The phosphorescent yield high-end proxy limits for these ROIs fall within the 10^{-2} to 10^{-1} range. There are even a few data points that approach or exceed 10^{-1} (R03_S10 & R20_S20) and such phosphorescence levels might be hard to believe, especially if the LED flash timing correction factor of $e^{+0.62}$ is also applied. One thing to keep in mind is that vampire pixels are known to bend drift field lines to produce regions with (apparently) $> 100\%$ QE. For example, the vampire pixel on R03_S10_C15 contains a group of pixels that can receive up to $15\times$ the target level in a flat exposure, so such large yields as we have seen here are perhaps not so mysterious after all.




Figure 187: Signal and wavelength response for phosphorescence expression (99% level) in ROIs of images for R01_S00. This is the prominent cosmetic seen in Fig. 173, which is apparently a vampire pixel.




Figure 188: Signal and wavelength response for phosphorescence expression (99% level) in ROIs of images for R02_S02. This is the diffuse phosphorescence that correlates with the coffee stains seen in Fig. 174. No extractions were performed on the vampire pixels found on the same sensor (R02_S02_C15 and R02_S02_C07).




Figure 189: Signal and wavelength response for phosphorescence expression (99% level) in ROIs of images for R02_S12. This is the structured phosphorescence that correlates with the coffee stains seen in Fig. 175.




Figure 190: Signal and wavelength response for phosphorescence expression (99% level) in ROIs of images for R03_S10. These describe regions including or near the bright/focusing vampire pixel seen in Figs. 176, 136b and 137b.




Figure 191: Signal and wavelength response for phosphorescence expression (99% level) in an ROI of images for R20_S20. These describe the prominent non-focusing vampire pixel seen in Figs. 136d and 137d.




Figure 192: Signal and wavelength response for phosphorescence expression (99% level) in ROIs of images for R43_S11. These describe bright, diffuse transient regions seen in Figs. 177 and 137f, which apparently turn off completely when the HV Bias is *off*.




Figure 193: Signal and wavelength response for phosphorescence expression (99% level) in ROIs of images for R43_S20. These include some of the the highly structured *snowflake-like* transient regions seen in Figs. 137h and 178.

H References

- Astier, P., Regnault, N., 2023, A&A, 670, A118 (arXiv:2301.03274), doi:10.1051/0004-6361/202245407, ADS Link
- Astier, P., Antilogus, P., Juramy, C., et al., 2019, A&A, 629, A36 (arXiv:1905.08677), doi:10.1051/0004-6361/201935508, ADS Link
- Banovetz, J., Utsumi, Y., Meyers, J., et al., 2024, arXiv e-prints, arXiv:2411.13386 (arXiv:2411.13386), doi:10.48550/arXiv.2411.13386, ADS Link
- [DMTN-276]**, Banovetz, J., Utsumi, Y., Slater, C., 2024, *Effects of Persistence on E2V Sensors and its Impacts on DC2 Data*, Data Management Technical Note DMTN-276, NSF-DOE Vera C. Rubin Observatory, URL <https://dmtn-276.lsst.io/>
- Doherty, P.E., Antilogus, P., Astier, P., et al., 2014, In: Holland, A.D., Beletic, J. (eds.) High Energy, Optical, and Infrared Detectors for Astronomy VI, vol. 9154 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 915418, doi:10.1117/12.2056733, ADS Link
- Esteves, J.H., Utsumi, Y., Snyder, A., et al., 2023, PASP, 135, 115003 (arXiv:2308.00919), doi:10.1088/1538-3873/ad0a73, ADS Link
- Janesick, J.R., 2001, *Scientific charge-coupled devices*, ADS Link
- Juramy, C., Antilogus, P., Le Guillou, L., Sepulveda, E., 2020, arXiv e-prints, arXiv:2002.09439 (arXiv:2002.09439), doi:10.48550/arXiv.2002.09439, ADS Link
- Lange, T., Nordby, M., Pollek, H., et al., 2024, In: Bryant, J.J., Motohara, K., Vernet, J.R.D. (eds.) Ground-based and Airborne Instrumentation for Astronomy X, vol. 13096, 130961O, International Society for Optics and Photonics, SPIE, URL <https://doi.org/10.1117/12.3019302>, doi:10.1117/12.3019302
- Park, H., Nomerotski, A., Tsybychev, D., 2017, Journal of Instrumentation, 12, doi:10.1088/1748-0221/12/05/C05015, ADS Link
- Park, H., Karpov, S., Nomerotski, A., Tsybychev, D., 2020, Journal of Astronomical Telescopes, Instruments, and Systems, 6, 011005, doi:10.1117/1.JATIS.6.1.011005, ADS Link

Polin, D., Snyder, A., Lage, C., Tyson, J.A., 2025, arXiv e-prints, arXiv:2502.05418 (arXiv:2502.05418), ADS Link

[Document-10512], Rasmussen, A., 2011, *LSST Dynamic Range and its Relation to Sensor Full Well Depth*, Informal Construction Document Document-10512, NSF-DOE Vera C. Rubin Observatory, URL <https://ls.st/Document-10512>

[LCA-19636], Rasmussen, A., 2024, *Focal Plane Fill Factor Verification Report*, Camera Controlled Document LCA-19636, NSF-DOE Vera C. Rubin Observatory, URL <https://ls.st/LCA-19636>

Roodman, A., Rasmussen, A., Bradshaw, A., et al., 2024, In: Bryant, J.J., Motohara, K., Vernet, J.R.D. (eds.) *Ground-based and Airborne Instrumentation for Astronomy X*, vol. 13096 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 130961S, doi:10.1117/12.3019698, ADS Link

SLAC National Accelerator Laboratory, NSF-DOE Vera C. Rubin Observatory, 2025, *The LSST Camera (LSSTCam)*, URL <https://www.osti.gov/servlets/purl/2571927>, doi:10.71929/rubin/2571927

Snyder, A., Longley, E., Lage, C., et al., 2021, *Journal of Astronomical Telescopes, Instruments, and Systems*, 7, 048002, doi:10.1117/1.JATIS.7.4.048002, ADS Link

Snyder, A., Polin, D., Lage, C., Tyson, J.A., 2024, In: Holland, A.D., Minoglou, K. (eds.) *X-Ray, Optical, and Infrared Detectors for Astronomy XI*, vol. 13103 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 1310321, doi:10.1117/12.3020540, ADS Link

[LCA-20583], Thayer, G., 2024, *Guider Requirements Verification Report*, Camera Controlled Document LCA-20583, NSF-DOE Vera C. Rubin Observatory, URL <https://ls.st/LCA-20583>

Utsumi, Y., Antilogus, P., Astier, P., et al., 2024, In: Holland, A.D., Minoglou, K. (eds.) *X-Ray, Optical, and Infrared Detectors for Astronomy XI*, vol. 13103 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 131030W, doi:10.1117/12.3019117, ADS Link

I Acronyms

Acronym	Description
1D	One-dimensional
2D	Two-dimensional
3D	Three-dimensional
AC	Access Control
ADC	atmospheric dispersion corrector
ADU	Analogue-to-Digital Unit
B	Byte (8 bit)
BOT	Bench for Optical Testing
CA	Certificate Authority
CAM	CAMera
CCD	Charge-Coupled Device
CCOB	Camera Calibration Optical Bench
CCS	Camera Control System
CMB	Cosmic Microwave Background
CMOS	complementary metal-oxide semiconductor
CSC	Commandable SAL Component
CTI	Charge Transfer Inefficiency
DAQ	Data Acquisition System
DC	Data Center
DDS	Data Distribution System
DOE	Department of Energy
E2V	Teledyne
EFD	Engineering and Facility Database
EO	Electro Optical
ESD	electrostatic discharge
FCS	Filter Changer System
FES	Filter Exchange System
FITS	Flexible Image Transport System
FRACAS	Failure Reporting Analysis and Corrective Action System
Gb	Gigabit
HV	High Voltage
ICD	Interface Control Document

IR	infrared
IR2	Interaction Region 2 (at SLAC)
ISR	Instrument Signal Removal
IT	Information Technology
ITL	Imaging Technology Laboratory (UA)
JSON	JavaScript Object Notation
L1	Lens 1
L2	Lens 2
L3	Lens 3
LATISS	LSST Atmospheric Transmission Imager and Slitless Spectrograph
LCA	Document handle LSST camera subsystem controlled documents
LED	Light-Emitting Diode
LOVE	LSST Operators Visualization Environment
LSST	Legacy Survey of Space and Time (formerly Large Synoptic Survey Telescope)
LSSTCam	LSST Science Camera
MB	MegaByte
MC	Monte-Carlo (simulation/process)
MCM	Master Control Module
MTCamera	Main Telescope Camera
Mb	Megabit (1000000 bit)
NFS	Network File System
OCPS	OCS Controlled Pipeline System
OCS	Observatory Control System
OpSim	Operations Simulation
PCS	Pumped Coolant System
PCTI	Parallel Charge Transfer Inefficiency
PD	Program Development
PS	Project Scientist
PSF	Point Spread Function
PTC	Photon Transfer Curve
QE	quantum efficiency
REB	Readout Electronics Board
ROI	Region of Interest

RTM	Raft Tower Module
S3	(Amazon) Simple Storage Service
S3DF	SLAC Shared Scientific Data Facility
SAL	Service Abstraction Layer
SCTI	Serial Charge Transfer Inefficiency
SLAC	SLAC National Accelerator Laboratory
TMA	Telescope Mount Assembly
UCD	Unified Content Descriptor (IVOA standard)
USA	United States of America
UT	Universal Time
UTC	Coordinated Universal Time
XML	eXtensible Markup Language