



Vera C. Rubin Observatory
Rubin Observatory Operations


The LSST Camera commissioning

LSST Camera Team, Pierre Antilogus, Alexandre Boucaud,
Claire Juramy-Gilles, Yousuke Utsumi, Sean Patrick MacBride,
Stuart Marshall, John Banovetz, and Theo Schutt

CTN-007

Latest Revision: 2025-12-12

DRAFT



Abstract

This technote describes what we found during the LSST Camera commissioning.

Draft

Change Record

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

Document source location: <https://github.com/lsst/ctn-007>

Draft

Contents

1 Introduction	1
2 Further analysis from LSST Camera Team et al. (SITCOMTN-148)	1
3 Camera operations	1
3.1 Recovery of CCDs	1
3.2 Cryo system	1
3.3 CCS development	10
3.4 Firmware development	10
3.5 Observation control	14
3.6 FES commissioning	14
3.7 Guider star finder	14
3.8 Historical telemetry of infrastructure cut events: electricity, glycol, dynalene and air	14
4 Camera optimization	14
4.1 Sequencer changes	15
4.1.1 Long range correlation	15
4.1.2 3sec sequencer file development	17
4.2 Re-implementation of "clear" operation during the idle time	19
4.3 Introduction	19
4.4 Background/Context for the sequencer, DAQ and CCS	19
4.5 IdleClearing	20
4.6 Science and Guider Details	21
5 Camera on-sky performance	22
6 Shutter	22
6.1 Shutter Timing issue	22
6.1.1 Overview of the shutter operation	23
6.1.2 Verification of the Observatory's PTP	25
6.1.3 Verification with the spare shutter system	26

6.1.4 Summary	29
6.2 Shutter blade trajectory fitting	30
7 Acknowledgements	31
8 References	31
9 Acronyms	32

Draft

The LSST Camera commissioning

1 Introduction

2 Further analysis from LSST Camera Team et al. (SITCOMTN-148)

3 Camera operations

3.1 Recovery of CCDs

3.2 Cryo system

- Temp sensitivity issue – we can put some episode and what we have done in this document, Theo will describe a few episodes

- Brief introduction of the refrigeration system
- Brief introduction of Cryo Temp sensitivity issue
- Detail of the cryo-temp sensitivity issue
- Detail of the Cryo 1 issue
- Detail of the Cryo 3 issue
- Analysis on the GC
- Analysis on the Oil
- Guide for adding refrigerants

Brief introduction of the refrigeration system

The refrigeration system for the LSST camera is composed of a cryogenic or “cryo” system to cool the focal plane CCDs, and an independent “cold” system to remove process heat from

the imaging sensor's readout-electronics boards (REBs) which are contained within the same vacuum vessel. The cryo system operates at approximately -130°C , removing 500W, and the cold system operates at approximately -50°C , removing 1100W. The two systems act independently, save for the radiative coupling of the two distinct evaporators within the cryostat vacuum vessel, and through weak radiative coupling and conduction paths in each Raft Tower Module (RTM). In addition to performance requirements on the evaporator temperature for nominal heat loads, its temporal stability, and spatial uniformity, there are several operational requirements that address the systems' performance under various extreme process load conditions, and a strict environmental interface (heat and vibration) to the telescope's dome enclosure, such as to avoid any impact to imaging.

The refrigeration system must meet the requirements shown in Table 1.1.1 which represent the interface between the telescope and the camera. These ensure that while operating, the refrigeration system does not impact the quality of the telescope images, neither by influencing the dome environment, nor by creating mechanical vibrations that could feed through to the focal plane. The refrigeration system must also operate in the anticipated temperature and humidity ranges at the observatory summit. These environmental requirements are shown in Table 1.1.2.

Finally, Tables 1.1.3 and 1.1.4 show the performance requirements for the cold and cryo systems, respectively. These include the temperature requirements for nominal and extreme process heat loads, as well as the temperature stability, uniformity and controllability of the two evaporators

Table 1.1.1: Summit Facility Interface between Telescope and Camera

Refrigerant transfer line skin temperature	Within $\pm 4/-15$ $^{\circ}\text{C}$ of ambient dome's temperature	Applies to all refrigerant line surfaces exposed to dome's ambient environment. Insulation is optional wherever possible.
Compressor cabinet outer skin temperature	Within $\pm 1/-3$ $^{\circ}\text{C}$ of dome's ambient temperature	Foam insulated cabinet walls, and coolant used to control the cabinet interior temperature

Table 1.1.2: Operating Dome Environment Requirements

Ambient working temperature range	-5 °C to 30 °C	System should deliver the required refrigeration capacity in the required evaporator temperature range.
Ambient turn-on temperature range	-10 °C to 30 °C	System should be able to turn-on and operate, but the refrigeration capacity may be reduced.
Ambient survival temperature range	-15 °C to -10 °C & 30 °C to 40 °C	System should survive such temperature extremes. The system does not need to be running in this condition. System should deliver the required refrigeration capacity in the required evaporator temperature range.
Ambient working humidity range	15% to 60%	System should survive these humidity extremes but does not need to be running in this condition. If the refrigeration system is running, the refrigeration capacity may be reduced.
Ambient survival humidity range	10% to 15% & 60% to 90%	

The cooling requirements for the LSST Camera are rather unique amongst astronomical instruments, driven by the large focal plane area with 189 distinct CCDs, the proximate location of the front-end electronics, and the parallel digitization and readout with large power dissipation, all occupying the same vacuum space. In addition, the cooling system must not influence the dome environment of the observatory itself. Heretofore, smaller but similar systems could use conventional coolants such as LN_2 . However, to provide adequate refrigeration capacity with such coolants would have required either a large refillable vacuum insulated reservoir directly behind the camera, or a set of flexible, large cross section vacuum insulated transfer lines running up to the camera from a reservoir and pump at the base. Both solutions would

require similar transfer lines to return cold boiled off N_2 vapor, the combination of which proves highly inconvenient. Such would also be the case with more modern cryocoolers (e.g.: Polycold) that locally produce cold refrigerant mixes, but which then must be transported to the evaporator and returned to the cryocooler.

Instead of these cumbersome conventional solutions, the design that has been selected incorporates two distinct but classic vapor compression refrigeration systems (“cold” and “cryo”). The cryo refrigeration system consists of 6 individual circuits, which utilizes the auto-cascade refrigeration technology. The refrigerants are compressed and cooled back to ambient dome temperature (by compressors and condensers remote from the camera), and then transported at high pressure through the dome, and up to the camera. There they are expanded to produce the required cooling of the evaporators. At the rear of the camera each system has counterflow heat exchangers which pre-cool these incoming refrigerants before expansion, utilizing the cold returning refrigerant. This refrigerant then exits the camera, now close to the dome temperature and at a low pressure and returns back to the remote compressors.

The cold system utilises conventional cascade refrigeration technology to cool the heat transfer fluid to -50 C at the compressor cabinet, the heat transfer fluid is pumped via a vacuum jacketed line to the camera, now at -48 C, to cool down the electronics. The heat transfer fluid then exits the camera, now at about -44 C, returns to the remote compressor cabinet via a vacuum jacketed line.

The transfer line system carries the refrigerants to the camera and returns them to the compressors – a total path length of >65 meters with a maximum elevation change of ~16 meters

Table 1.1.3: Cold Refrigeration System Performance Requirements

Temperature range for operation	-40 °C to -30 °C when the load ranges from 754 W to 843 W	Nominal load is 800W at -35 °C
Maximum survival load	1056 W	The system does not run away. The system does not run away. Trim heat and/or compressor operating frequency and /or hot gas bypass can be used when running the system with the minimum survival load.
Minimum survival load	490 W	Temperature across the cold plate
Temperature uniformity	+/-1.5 °C	The controller should adjust the heat load of the trim heaters to compensate the change of the process load when turning on/off electronics.
Temperature controllability	+/-2 °C when process load changes by 89 W	
Temperature drift	+/-1 °C at constant load	Duration not specified.

Table 1.1.4: Cryo Refrigeration System Performance Requirements

Cryo plate temperature range	-130 °C to -125 °C with load from 361 W to 498 W	Any temperature from -130 °C to -125 °C is accepted when the load is in this range. The nominal temperature and load are -127.5 °C @ 429.5 W.
Cryo plate maximum survival load	520 W	Cryo system should not run away. Use trim heat to increase the minimum survival load.
Cryo plate minimum survival load	219 W	Cryo system should not run away.
Cryo plate temperature uniformity	+/-1.5 °C +/-1.5 °C when process load changes	Temperature across the cryo plate The controller should adjust the heat load from the trim heaters to compensate the change in load from when turning on/off electronics.
Cryo plate temperature controllability	136 W +/-0.5 °C	
Cryo plate temperature drift	at constant load	Duration not specified.

Architecture and Thermal Design of the Cryostat

The cryostat is designed to be a vacuum insulating vessel ($<10^{-7}$ Torr) containing the cold and cryo systems' evaporators, the "coldplate" and the "cryoplate". Figure 1.3.1 shows a cross section of the cryostat with the principal thermal regions and typical operating temperatures. Refrigerants enter the cryostat, pass through the evaporators and return through 4.5" vacuum conflat flanges on the rear annular feedthrough plate. The refrigerant tubing itself is continuous and joint-free within the cryostat vacuum to eliminate any potential for a leak.

The refrigeration system components outside the cryostat (in the Utility Trunk) have an independent insulating vacuum, referred to as the heat exchanger (“HeX”) vacuum). Ceramic insulators are used throughout the system to isolate the electrical ground of the evaporators from the rest of the cryostat. Figure 1.3.2 shows the detailed routing and makeup of the system at the interface between the cryostat and the heat exchanger (HeX) vacuum system, including the mechanism for isolating grounds.

The cryoplate layout is shown in Figure 1.3.3. It is composed of a single machined and polished copper plate which is brazed to a stainless-steel plate to rigidize it, with cutouts for each of the Raft Tower Modules (RTMs). There are 21 “bays” where the science rafts are attached, and 4 corner bays for the corner rafts containing guider CCDs and wavefront sensors. The six stainless steel refrigerant lines – one for each cryo circuit - run continuously through the top and bottom of the six copper ribs, having been brazed in place. There are three lines on the top, and three on the bottom, and the refrigerant direction is organized to be counterflow. The rafts are bolted down to the copper on the front side of the orthogonal ribs.




Figure 1.3.1 Cross section of the cryostat with its principal thermal regions and typical design temperatures during operation (in red).




Figure 1.3.2 Detailed routing of the refrigeration system at the interface between the cryostat and the heat exchanger vacuum system.




Figure 1.3.3 Cryoplate (a) Routing of six refrigerant circuits in the copper cryoplate (b) with stainless steel stiffening frame attached.

Architecture of the Cryogenic Systems

The physical layout of each cryo system is shown schematically in Figure 1.4.1.1. With the exception of the transfer line lengths, a largely identical system supports the camera in the clean room Maintenance Facility at the Vera C. Rubin Observa-

tory.

All of the ground-based components reside in two temperature-controlled housings [“cryo cabinets”] attached to the underside of the Telescope Mount Assembly (TMA). The transfer lines to and from the camera (a liquid supply, a vapor supply, and a suction-return for each circuit) originate and return to these cabinets. They are a mix of orbitally welded - electropolished 316-L stainless steel tubing and flexible refrigerant lines which are insulated wherever possible. The counterflow HeX systems reside in a separate vacuum insulated housing within the camera’s Utility Trunk, just behind the cryostat

3.3 CCS development

- CCS

3.4 Firmware development

Science REB (REB)

Table 5: Science REB Firmware Version History

Version	Release Date	Key Modifications and Notes
0x313B5019	December 2025 (6.4ns: tested briefly on camera; 10ns: tested on ts8)	• Code Refactor Completion: Allows system clock selection on a per-target basis using common functional blocks. → The same VHDL code is used in both cases with only a parameter change. • Expected to definitively fix the ID PROM read instability issue (pending observation). • These libraries are now reusable in GREB/WREB firmware; remaining work is to apply them to devices unique to those boards.

Table 5 – Continued

Version	Release Date	Key Modifications and Notes
0x313a5015 - 0x313B5018	November 2025 (Not installed on camera)	• A series of releases used on ts8 for the development and testing of the 6.4ns firmware version. • All had various issues, but each was separately versioned before installation on ts8.
0x313a5014	September 2025 (Installed on camera: November 2025?)	• Functional Equivalence: This is the production version that ended up functionally equivalent to 0x31395012. → 0x3139500e with new functions: → – Ability to build firmware with new bias protection thresholds and sequencer overrides. → – Stability of CCD power across REB resets (rms_reset). • ID PROM Instability: Still suffers from the latent “ID prom instability,” a long-standing issue for which the root cause was never fully identified, despite many versions appearing stable.
0x313a5013	August 2025 (Briefly installed on camera)	• Nominally intended to be functionally equivalent to 0x31395012 under the hood. • Technical Stack Update: Migrated to the currently supported SLAC TID-ID firmware build system, SURF libraries, and Vivado 2025.1 compiler. • Code Refactoring: Shared code blocks moved to a dedicated repository (https://github.com/lst-camera-electronics/lst_reb). • Intended as preparation to support both 10ns and 6.4ns clocks. • Major Defect: Introduced a major bug resulting in maximum heater power being applied due to a register map conflict.

Table 5 – Continued

Version	Release Date	Key Modifications and Notes
		→ This bug was actually introduced in 0x31395011 but went unnoticed.
0x31395012	July 2025 (Not installed on camera)	Made Bias VUT and HV switch stable across REB resets.
0x31395011	May 2025 (Installed on camera R30/Reb1 only)	- Made CCD bias protection thresholds compile-time configurable. - Added a “sequencer override register” to fix sensor sequencer signals.
0x3139510F	April 2023 (Not installed on camera)	Included fixes discovered during Guider development, only relevant to the multi-sequencer version.
0x313950FC - 0x313950FF	(Various Release Dates) (Used on ts8 only)	A series of 6.4ns test versions built by naively modifying 0x3139500e.
0x3139500e	July 2022 (Production until November 2022)	Baseline Production Release.

Wavefront REB (WREB)

Table 6: Wavefront REB Firmware Version History

Version	Release Date	Key Modifications and Notes
0x1139400B	September 2025	• Updates: Migrated to the current SLAC TID-ID build system, SURF libraries, and Vivado 2025.1. • Code Refactoring: Shared code blocks moved to dedicated repository (https://github.com/lst-camera-electronics/lst_reb). • Intended as preparation to support both 10ns and 6.4ns clocks.
0x1139400A	July 2025	• Ensured stability of Bias VUT and HV bias switch across REB resets.
0x11394009	May 2025	• Added compile-time specification for bias protection thresholds (unused in WREBs).
0x11394008	April 2023	• Baseline Release Version.

Guider REB (GREB)

Table 7: Guider REB Firmware Version History

Version	Release Date	Key Modifications and Notes
0x2139210D	January 2025	• First production release of 2-sequencer firmware for Guider operation.
0x21392x0c	(Not Released)	• Development Dead End.
0x21392x0B	April 2023	• Versions where x=1 are 2-sequencer firmware; x=0 are single sequencer (full frame mode).

3.5 Observation control

- OCS integration kfanning

3.6 FES commissioning

- FES commissioning performance – Pierre Antilogus, Alexandre etc...

3.7 Guider star finder

@Aaron Roodman star finder in guider

This section describes the filter exchange system performance through commissioning.

3.8 Historical telemetry of infrastructure cut events: electricity, glycol, dynalene and air

@Homer Neal telemetries of infrastructure cut events, electricity, glycol, dynalene, and air

4 Camera optimization

4.1 Sequencer changes

This section describes the history of the sequencer changes from Run 7. Run 7 results in the v30 sequencer file. The changes made is documented in Section 5.2 in LSST Camera Team et al. (SITCOMTN-148). Since Run 7 was concluded, the following optimization targets are defined.

- long range correlation
- switching to the 3 sec readout timing

There was a schedule pressure to complete the sequencer optimization because the beginning of the Science Verification Survey was due, which required to simplify our testing.

4.1.1 Long range correlation

The measurement of brighter-fatter coefficients in LSST Camera Team et al. (SITCOMTN-148) had already shown an anomalously high value of the a_{00} coefficient for one specific e2v sensor: R23_S02 (see Figure 34 in that note). This measurement is derived from fits over the variance and spatial covariance curves versus the flux. The curves for the channels in that particular sensor were in fact affected by an excess of variance, which affected not only the noise measurement in bias frames, but also introduced an extra flux-dependent component in the variance. This degraded the gain measurements from the PTC (using C_{00}/μ). This issue also manifested as a flux-dependent covariance between amplifiers in the same sensor, and as an excess in the correlation with the next pixel in the serial direction (a_{10}). One way to detect this issue in a systematic manner is to compute the slopes of the covariances between channels versus the flux, and average them over all channels in the sensor (as done in Figure 1). Using this metric, other sensors were found to exhibit a similar issue, but to a smaller degree (same Figure, left plot). Other related metrics (noise correlation between channels, correlations between channels) also pointed to the same sensors.

Since the issue affected some sensors uniformly, and was exhibiting flux-dependency, investigations focused on the CCD readout clocks, and in particular on the pixel reset clock (RG). To explain a potential variation from sensor to sensor, it should be noted that while the REB FPGA drives all three sensors with the same readout sequence, the actual clocks are generated on a per-sensor basis. While the jitter on the FPGA timing signal is in the range of 100 ps,




FIGURE 1: Averages of the slopes of the covariances between channels, for each sensor. Left: during test run 7. Right: on sky, with the modified 'RGn20' readout sequence.

the jitter in the chain that translates that signal into a clock can be 1-3 ns. In addition, the RG clock is faster than the serial clocks, due to a lower capacitive load.

For sensors like R23_S02, plots acquired in scan mode at the raft level showed a strong variability (dispersion over consecutive scans) in the timing of the clock transitions around the time that the pixel reset (RG) clock was lowered. In the e2v sequence for pixel readout, this drop of the RG clock was timed simultaneously with the drop of the S2 clock (which transfers charges to the S3 phase, in preparation for transfer to the readout node). The same scans also showed a much stronger variability in the level of the signal after the pixel reset, up to 1000 electrons instead of the expected 50-100 electrons from kTC noise. Even with the Dual Slope Integration that is performed in the CCD readout chain, such high variability would explain an extra component in the noise.

We then tested a series of modifications to the timing of the readout sequence. In a sensor that did not exhibit the extra covariance at the nominal timing, we were able to induce it by triggering the S2 clock drop 10 ns before the RG drop. Conversely, tests on sensors that exhibited the issue showed that delaying the S2 drop by an increment of 10 or 20 ns after the RG drop removed the issue (Figure 1, on the right). This 20 ns delay was selected as the safer option, and added to the e2v readout sequence (tagged as v30_RGn20, then adopted as v31). A similar modification was tested on ITL sensors, but there was no effect on sensors that exhibited correlated noise.

v32 has a stuff related to idle clear, no functional changes in other parts v33 is reordered in

a one-part sequencer in timing. Whether inserting a 20ns RG before (as BufferRGp) or after the BufferS+TimeS (as BufferRGn) time slice. No change in the total timing.

4.1.2 3sec sequencer file development

The observing strategy has evolved from two snaps of 15 sec exposures to a single 30 sec exposure per visit. By eliminating the intermediate readout, the overall timing budget is relaxed: one full readout cycle no longer needs to be performed during a visit. This reduction frees up readout time that can be reallocated to lower readout noise. Consequently, the development of a 3 s sequencer (instead of the existing 2sec) was initiated to take advantage of this additional margin.

We examined the times slices in the sequencer files and identified four key parameters for optimization:


- ISO1 – Time for ASPIC clamp/reset to start of the first integration to measure the baseline (Ramp Down; RD)
- ISO2 – Time between S3 down and start of ASPIC Ramp Up (RU)
- RampTime – the integration time of a pixel read.
- TimeP – the unit element of the Parallel Transfer (parallel timing).
- BufferP - The time period before each parallel transfer starts.

Previous studies (Utsumi et al., 2024; Snyder et al., 2024) examined how variations in ISO1, ISO2 and RampTime affect read-out noise; they also demonstrated that altering Parallel Timing (TimeP) can increase full-well capacity. ISO2 is additionally known to reduce crosstalk. For the v30 firmware, BufferP is set to 3360ns, raising the question of whether the entire parallel clocking should be shifted forward as in the original vendor sequence. In fact this change Because extending ISO1 degraded performance on some E2V sensors, our investigation was confined to the following scenarios:

1. **Longer integration time per pixel (3s-v0 RT)** – increases from 380 to 620 (1.6)
2. **Extended ISO2 duration (3s-v0 ISO2)** – lengthens the readout interval for each row.

3. **Longer parallel transfer timing (3s-v0 Parallel)** – allows more time for charge transfer between columns.

Design constraints emphasized minimal deviation from the baseline sequencer, preserving parameters that had been previously tuned.



Sequencer files for each option were produced and peer-reviewed. Validation involved bias checks, TS8 hardware testing for ITL, and UCD testing for E2V.

Noise measurements indicated that the ISO2 extension (option 2) delivers performance on par with the baseline while achieving the targeted 3-second readout time. Consequently, option 2 was adopted as the new production sequencer (3s-v1).

The total system gain is determined by three parameters: ramp time, ASPIC gain, and ASPIC RC. Since changing the Ramp Time scales the gain roughly linearly we also need to adjust the ASPIC gain to compensate the gain change from the Ramp Time in order to achieve the similar level of the output. The similar output signal was needed not to trigger so-called bias shift for the high ADC output.

The ASPIC RC parameter changes the selectable resistor in the integrating amplifier stage of the ASPIC, changing its RC time constant. The compensation can be achieved by setting RC from 14 to 2 to give the ratio of 1.6.

- E7662 3s RT sequencers: https://s3df.slac.stanford.edu/data/rubin/lstcam/E7662/w_2025_26/
- E7697 3s Parallel sequencers: https://s3df.slac.stanford.edu/data/rubin/lstcam/E7697/w_2025_26/
- E7714 3s ISO2 sequencers: https://s3df.slac.stanford.edu/data/rubin/lstcam/E7714/w_2025_26/

<https://rubin-obs.slack.com/archives/C07QJMPFLQ2/p1750880738722919>

- R30/S12 and R23/Reb10 – I will try to collect all the information from Slack

4.2 Re-implementation of “clear” operation during the idle time

This section describes the implementation of the “idleClear” mechanism used on the LSSTCam to flush charges when the CCDs are otherwise idle.

4.3 Introduction

In addition to intentional exposures, the LSSTCam sensors accumulate signal when they are idle, via light leaks, time dependent defects, and darkcurrent. When the sensors are idle, the power consumption changes affect the sensor and electronics temperatures. To keep the sensors ready for exposure and the signal chain at a steady temperature, we require a mechanism which clears the charges semi-continuously at a constant power level comensurate with normal survey observing. This mode is named “idleClearing”.

4.4 Background/Context for the sequencer, DAQ and CCS

The core LSSTCam readout is implemented in an FPGA via a firmware engine dubbed “the sequencer” which emits a sequence of values where each of 14 significant bits drive one of the control signals for the CCD pins, dual slope integration, digitization and data transfer. Each FPGA is integrated into a Raft Electronics Board (REB) with the electronics for up to 3 CCDs and resides with the CCDs inside the LSSTCam vacuum chamber. The sequencer is loaded with a set of “sequencer programs” that define the sequence of output states necessary to

accomplish CCD readout. This implementation architecture is limited in that each defined “sequencer program” is triggered, runs to completion and ends in an idle state. The sequencer does not support running a program as an interruptible infinite loop. As a result, the mechanism for idleClearing must be implemented outside of the FPGA and hence at the CCS or DAQ level. However, the sequencer is suitable for taking survey image data where each image is the result of one triggered run of a sequencer main program which reads all or a portion of the CCD. These sequencer programs are parametrized to allow dynamically adjusting readout for overscan, ROIs (region of interest) as one might expect. The sequencer programs are loaded into the FPGA and triggered by or through interfaces provided by the DAQ (Data Acquisition) system (described elsewhere).

The complete LSSTCam readout comprises two distinct sets of sequencers: those that support the 8 guider CCDs and those that support the 189 science and 4 wavefront CCDs. In the DAQ nomenclature, each set is referred to as a “partition” of the complete DAQ. Each partition is operated synchronously by a distributed trigger generated in the DAQ. For science/wavefront sensors, the trigger is initiated by the CCS (via the DAQ interface), once per image at the end of an exposure to read out in 3 seconds. During science sensor readout, the guider CCDs are intentionally static. During science sensor charge integration (exposure) the guider CCD readout is periodically triggered (nominally 9 Hz), using an ROI (region of interest). The guider readouts are triggered by processes internal to the DAQ using ROI parameters requested by the CCS. When no science sensor exposure or readouts are occurring, all CCDs would be “idle”. This is the issue to be mitigated by idleClearing.

4.5 IdleClearing

It is easiest to describe idleClearing in terms of one segment of one CCD. Since all of the segments for all CCDs in a partition are triggered synchronously one segment is sufficient. For this discussion, a single segment has N rows and M columns where N, M include any overscan. The standard values of (N, M) are $(2048, 576)$. To match the goal of constant power, we intend to read out pixels at the same average rate used during survey operation but we do not transfer any data outside of the REB. We define a period P , and a number of rows n . We set the sequencer parameters for an ROI of nominal width but with only $n < N$ rows. Then we arrange to trigger the sequencer every P seconds. This yields a readout rate of $(n \cdot M / P)$ pixels/sec to be compared with $(N \cdot M / (3 + \text{exptime}))$ which is 35000 for the nominal 31 second exptime.

4.6 Science and Guider Details

For the science partition we are choosing $n:100$, $P:1.6$ so using $M=576$ yields 36000 pix/sec. The plots for TotalRebPower during idle periods and observing match up as shown. [insert plot or two here with some explanation]

For the guider partition, we are using $n:10$, $P:0.090$ and $M=576$ to yield 64000. We compare to the rate during nominal guiding where we read a 50×50 ROI at 9Hz. Ignoring the flushed pixels this gives 22500 pixels/sec. However this needs a bit more thought to account for the power used in shifting rows and columns at maximum rates to access the ROI which also consumes quite a bit of power. [This needs a bit more flesh, but at first glance we should reduce P a bit if we need to match.] [insert plot or two here with some explanation]

With power matched and charges cleared, there is still the issue of latency when the system needs to be used for the intended purpose of taking images or guider stamps. In this case, the loop must be stopped and the system must transition into a state ready to operate. The science and guider cases are slightly different.

Considering the science partition case first where the loop is implemented as a periodic task in java. The task is scheduled periodically so it just triggers (via the DAQ interface) the sequencer to run. The period is chosen to be longer than one triggered sequencer execution. When it is time to stop idleClearing the task is scheduled task is cancelled and the system waits for confirmation that any in progress triggered run of the sequencer is completed. Hence an upper bound on the time to stop is the time for one triggered run plus a very small amount of overhead. In the case of $n:100$ rows this amounts to 150 milliseconds. This generally will not delay any images during survey imaging.

In the guider partition, the loop is implemented on an embedded linux node where it would not be appropriate to run java. The main loop has two tasks in series: first the sequencer is triggers (reading out 10 rows) and then the system sleeps for 80ms. Neither task is interruptable but the series can be interrupted at the end of the reading or at the end of the sleep. Since the reading of 10 rows takes 12 ms, the upper bound on latency to stop the idle clearing in the guider is 80ms. As with the science partition, the latency to stop idleClearing in the guider generally does not delay transitioning from idleClearing to survey operations.

Coordination and side-effects: managing the sequencer

Add a discussion of the issue that the sequencer is a single trigger for each partition and when running, has interactions with other aspects of REB/CCD management. So for example we would not want to turn on or off the CCDs with the sequencer running. Similarly if the focal-plane process for idleClear is periodically triggering the trigger for the science partition and we attempt to power cycle a REB, when the reb reboots it may/will receive a trigger before the sequencer is fully loaded in the REB registers. This is a situation we don't want to happen. More on this topic....

- eoPipe – Jim
- Correlated noise – Pierre Astier or Pierre Antilogus
- PTC – Pierre Astier

5 Camera on-sky performance

- Bright star test – SB
- On-sky persistence test – just showing a dark image after a bright star to say no persistence
- Pinhole repeatability tests – YU
- Phosphorescence study with u and g band exposures – SB
- persistency – YU

6 Shutter

6.1 Shutter Timing issue

Since the first photon observation was made, the shutter system has exhibited issues including the blade overrun, leaving the shutter partially opened, and discontinuities in motion profiles (Utsumi et al., CTN-003). This section presents the further analysis and results of follow-up tests, and current understanding and remaining open questions.

6.1.1 Overview of the shutter operation

The camera shutter is controlled by the Beckhoff TwinCAT PLC system. The CCS sends commands to a Beckhoff CX5120 controller, which is an industrial PC running Windows 7 Embedded Standard along with a real-time kernel called TwinCAT 3. This is the gateway and manages the PLC system interconnected by EtherCAT. The PLC system consists of various "terminal" modules that can take different roles, such as controlling digital and analog I/O as well as the timing module, so-called EL6688. EL6688 provides an external reference to synchronize the clock among terminals with the GPS clock in a better accuracy than $1\mu\text{s}$.

In this system, there are concepts of 4 different clocks:

1. The CCS clock synchronized by NTP, which is based on UTC,
2. The TwinCAT clock, which is supposed to be synchronized by the host Windows when it is booted
3. Distributed clock (DC) in the TwinCAT PLC system. When the shutter blade passes the hole sensors placed along with the shutter rail, the PLC is programmed to trigger the timing measurement based on DC. This timing information is the foundation of the motion profile as well as the entire timing control of the PLC system.
4. External reference clock from EL6688. The time source is a GPS clock in the observatory and the synchronization is provided by Precision Timing Protocol, which is expected to provide sub-microsecond timing accuracy through network delay measurements.

We investigate the consistency in the timings from logging messages retroactively. The focus is given on the CCS clock and the DC time reported by the motion profile. The UTC timestamps in the MCM log ('sent takeExposure to shutter1") were paired with TAI timestamps from motion profiles ('startTime.tai.isot"). Because TAI leads UTC by 37 seconds, the expected difference is approximately: *37 seconds + system/network jitter*. Plotting these deltas revealed three operational periods in Figure 2. There are three distinct periods in this plot.

April-June 2: The EL6688 was intermittently synchronized (the reported PTP state was unstable, bouncing between LISTENING (4) and SLAVE (9) states). However, evidence suggests that DC time was free-running. The shutter became unstable near the end of this period, including discontinuities in the motion profiles and blade overrun issues. This period ended after power




FIGURE 2: Timing between the CCS time when takeExposure was recorded and the time when the motion profile records the startTime. This doesn't reflect the actual shutter motion timing.

was initialized to the shutter system. The underlying cause remains unknown. A reduction in TCPSynRetrans was observed, which may indicate a Linux kernel-level anomaly.

June 3–June 7: The EL6688 synchronized to a MOXA boundary clock, but synchronization quality was poor. Observed scatter reached 50 ms. The MOXA switch is an internal L2 network switch inside the Camera (EDS-G516E-T). As described in Utsumi et al. (CTN-003), the PTP functionality of the MOXA switch does not appear to function correctly with the observatory network. The communication with the vendor confirmed that this device only supports the software-based PTP and does not support the hardware-based PTP which can explain the poor quality of the synchronization.

June 8–Dec 2 (present): The EL6688 synchronized to an observatory boundary clock on a leaf switch (10.17.200.33) after three hops from the Grand Master Clock, where the GPS clock is connected. Timing has remained stable.

6.1.2 Verification of the Observatory's PTP

Further analysis was conducted using the spare system lsstcam-shutter02 in the clean room. This system is connected to the observatory's network. This network is specially configured to pass the PTP packets. The network routes through an identical MOXA switch, providing an opportunity to investigate whether the observatory's PTP system functions as intended and to identify any issues with the MOXA switch. Tests conducted in the clean room revealed successful PTP synchronization using a MOXA operating as an L2 switch by the following command:

```
[yutsumi@lsstcam-shutter02 ~]$ sudo ptp4l -f /etc/ptp4l.conf -i enp0s31f6 -m
[sudo] password for yutsumi:
ptp4l[503690.709]: selected /dev/ptp1 as PTP clock
ptp4l[503690.710]: port 1 (enp0s31f6): INITIALIZING to LISTENING on INIT_COMPLETE
ptp4l[503690.710]: port 0 (/var/run/ptp4l): INITIALIZING to LISTENING on
    ↪ INIT_COMPLETE
ptp4l[503690.710]: port 0 (/var/run/ptp4lro): INITIALIZING to LISTENING on
    ↪ INIT_COMPLETE
ptp4l[503696.415]: port 1 (enp0s31f6): new foreign master fc59c0.ffff.214803-29
ptp4l[503712.418]: selected best master clock 000efe.ffff.011598
ptp4l[503712.418]: port 1 (enp0s31f6): LISTENING to UNCALIBRATED on RS_SLAVE
```

```
ptp4l[503714.237]: port 1 (enp0s31f6): minimum delay request interval 2^5
ptp4l[503714.620]: master offset -311009524 s0 freq +6354 path delay 10727
ptp4l[503715.623]: master offset -311010780 s1 freq +5099 path delay 10727
ptp4l[503716.621]: master offset 474 s2 freq +5573 path delay 10727
ptp4l[503716.621]: port 1 (enp0s31f6): UNCALIBRATED to SLAVE on MASTER_CLOCK_SELECTED
ptp4l[503717.623]: master offset 1484 s2 freq +6725 path delay 10727
ptp4l[503718.625]: master offset 257 s2 freq +5943 path delay 10727
ptp4l[503719.623]: master offset -26 s2 freq +5737 path delay 10727
ptp4l[503720.628]: master offset -184 s2 freq +5571 path delay 10727
ptp4l[503721.626]: master offset 8 s2 freq +5708 path delay 10727
ptp4l[503722.626]: master offset 443 s2 freq +6145 path delay 10727
ptp4l[503723.631]: master offset 83 s2 freq +5918 path delay 10727
ptp4l[503724.631]: master offset -434 s2 freq +5426 path delay 10727
ptp4l[503725.635]: master offset -149 s2 freq +5581 path delay 10727
```

The offset is in nanoseconds unit and s2 indicates the servo is in the locked state. As indicated in the output, the offset converges quickly. The key in the configuration is that the priorities have to be 255 since the priority of the grandmaster is set to be 255, otherwise lsstcam-shutter takes the grandmaster role in the PTP system which screws up the PTP synchronization. These tests confirm the observatory's PTP functionality.

There is another machine in the clean room Dell AIO, aio06. This machine was also used to test the PTP. The hardware clock behaved poorly under frequency adjustment, but acceptable synchronization was achieved using weakened servo parameters:

```
[lsstcam-aio06] % sudo ptp4l -f /etc/ptp4l.conf -i enp0s31f6 -m -H --step_threshold=1
➡ --pi_proportional_const=0.01 --pi_integral_const=0.0001
```

6.1.3 Verification with the spare shutter system

Since the PTP synchronization has been successfully established with linux computers through the MOXA switch, the focus moves onto the actual PTP terminal module EL6688. EL6688 was particularly difficult because the interface only supports 100Mbps connection. There was a doubt whether the link speed conversion does not interfere with the PTP synchronization. We used several tools to investigate the PTP synchronization comprehensively.

'readPtpDiag.sh' is the tool to communicate with Beckhoff controller over the network, which provides a partial view of the variables in the PLC. The following example is taken from Utsumi et al. (CTN-003).

```
[ccs@lsstcam-shutter02 ~]$ /usr/local/bin/readPtpDiag.sh
HCU date and time: 2025-06-09 17:58:35.614918
PTP version: PTPv2 (32)
PTP state: SLAVE (9)
Clock ID (hex): 00 01 05 ff fe 3e 4c 43
Master clock ID (hex): fc 59 c0 ff ff 21 48 03 00 1d
Grandmaster clock ID (hex): 00 0e fe ff fe 01 15 98
Offset from master (ns): -11,662
Mean path delay (ns): 122,992
Steps removed: 4
Sync msg sequence no.: 37388
Time scale: PTP/TAI (1)
Offset from UTC (sec): 37
UTC offset valid? True
Leap61? False
Leap59? False
Epoch no.: 0
External (PTP) time: Wed, 16 Jul 2025 17:40:44.907797 UTC
```

The output has a few distinct states every power cycle. "Offset from Master" is sometimes quite large, as large as 10^9 nanoseconds (or 1 second), and sometimes it is small as an order of $10\mu\text{s}$ or could be smaller, regardless of the "PTP state". And sometimes the PTP state never becomes SLAVE(9) but LISTENING(4).

Windows's program TwinCAT XAE shell (TcXaeShell) provides a live and full view of the variables in the PLC system. Figure 3 shows an example screenshot of TcXaeShell when the vendor-provided example project is running. DcToTcTimeOffset and DcToExtTimeOffset are one of the important variables because we found that they tell which of either TwinCAT or PTP clock is used to synchronize to DC. It appears that only one of them becomes non-zero, and the other becomes zero but randomly. Whether which clock is used is not deterministic as long as we use TwinCAT version 3.14022.36.

bSynchronized and nTimeDiff are also important parameters. These are unique in the vendor-provided project. These tell us whether EL6688 is synchronized with PTP or not and how much delay is observed. The meaning of the synchronization in the vendor implementation is broad. The synchronization is established regardless of whether the DC is synchronized to the Beckhoff Windows or EL6688, since EL6688 gives a reference clock and the actual clock synchronization is done in the Beckhoff Windows system using the built-in functions. nTimeDiff seems to be different from Offset from master.




FIGURE 3: A screenshot of the vendor provided sample program. DcToTcTimeOffset, DcToExtTimeOffset and bSynchronized as well as nTimeDiff are displayed in this screenshot.

Another diagnostic tools is a generic open-source software, tcpdump.

```
sudo tcpdump -i enp0s31f6 -n -s 0 'ether proto 0x88f7 or udp port 319 or udp port 320' -a
```

The example output of tcpdump is the following, focusing on a communication pair between the boundary clock and the client (EL6688) when they exchange "delay request" and "delay response".

```
12:33:01.777635 IP 139.229.175.116.ptp-event > 224.0.1.129.ptp-event: PTPv2, v1
  ↳ compat : no, msg type : delay req msg, length : 44, domain : 0, reserved1 : 0,
```

```

    ↪ Flags [none], NS correction : 0, sub NS correction : 0, reserved2 : 0, clock
    ↪ identity : 0x105fffe3e4c43, port id : 1, seq id : 8038, control : 1 (delay req
    ↪ msg), log message interval : 5, originTimeStamp : 1764419582 seconds, 136332373
    ↪ nanoseconds
12:33:01.777892 IP 10.17.200.33.ptp-general > 224.0.1.129.ptp-general: PTPv2, v1
    ↪ compat : no, msg type : delay resp msg, length : 54, domain : 0, reserved1 : 0,
    ↪ Flags [none], NS correction : 0, sub NS correction : 0, reserved2 : 0, clock
    ↪ identity : 0xfc59c0ffff214803, port id : 29, seq id : 8038, control : 3 (peer
    ↪ delay resp msg), log message interval : 5, receiveTimeStamp : 1764419618
    ↪ seconds, 777492282 nanoseconds, port identity : 0x105fffe3e4c43, port id : 1

```

This example confirms that EL6688 communicates with the boundary clock (a leaf switch) directly, regardless of the presence of the MOXA L2 switch. This message exchange is a part of the foundation to measure the delay in the communication, which enables precise timing synchronization as defined in the PTP specification. In this case, the time difference from originTimeStamp to receiveTimeStamp is 36.641159909 sec. Repeating observations of these packets found that delays are consistent over multiple days with a better accuracy than 1 μ s.

Although the difference is consistent, the difference itself is quite large. This large difference was observed when the large Offset from master was observed (DC is synchronized with the TwinCAT clock). On the other hand, if DC is synchronized with EL6688/PTP, the difference was quite small. This observation suggests that when the DC synchronized with the TwinCAT clock and the DC tries to keep the pace as referenced from EL6688 (PTP) but likely the absolute clock could be wrong if TWInCAT clock is not reasonably synchronized with NTP.

6.1.4 Summary

- The observatory PTP system provides a reasonable PTP synchronization.
- MOXA switch does not support the hardware clock, resulting in the poor PTP synchronization.
- Even EL6688 supports up to 100Mbps, reasonable synchronization can be established.
- TwinCAT's behavior whether DC is synchronized to TwinCAT controller or EL6688 is not

consistent and not configurable. Power cycle until it has a small enough offset from Master is needed operationally.

- **We also note that motion profile timing prior to June 7 should be considered unreliable; afterward, it appears trustworthy.**
- The jumps in the motion profile can be explained by intermittent adjustment based on the EL6688 to the TwinCAT system appears to play a role here.

Recommendations

- Correct the Puppet-managed chrony.conf on shutter01/shutter02 (LSSTCAM-146).
- Perform the Chris Walter timing test to interdependently check the timing accuracy. Walter (SITCOMTN-136)
- Consider to increase the proprieties (lower the number) of the observatory grandmaster/boundary clocks. Any misconfigured computer can mess up the PTP network as ptp4l's default is 128.

Open Questions

- Beckhoff's statement that EL6688 "does not adjust the distributed clock" remains unclear, especially regarding NTP/PTP coexistence.
- Compare MCM command timestamps with motion profile timestamps for deeper latency characterization.
- Determine whether shutter02 can store/export motion profiles for analysis.

6.2 Shutter blade trajectory fitting

@andy rasmussen shutter blade trajectory fitting

7 Acknowledgements

This material is based upon work supported in part by the National Science Foundation through Cooperative Agreements AST-1258333 and AST-2241526 and Cooperative Support Agreements AST-1202910 and AST-2211468 managed by the Association of Universities for Research in Astronomy (AURA), and the Department of Energy under Contract No. DE-AC02-76SF00515 with the SLAC National Accelerator Laboratory managed by Stanford University. Additional Rubin Observatory funding comes from private donations, grants to universities, and in-kind support from LSST-DA Institutional Members.

8 References

- [SITCOMTN-148]**, LSST Camera Team, Antilogus, P., Astier, P., et al., 2025, *LSST Camera Electro-Optical Test (Run 7) Results*, Commissioning Technical Note SITCOMTN-148, NSF-DOE Vera C. Rubin Observatory, URL <https://sitcomtn-148.lsst.io/>, doi:10.71929/rubin/2583999
- 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
- 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, 131030W, International Society for Optics and Photonics, SPIE, URL <https://doi.org/10.1117/12.3019117>, doi:10.1117/12.3019117
- [CTN-003]**, Utsumi, Y., Schutt, T., Liang, S., et al., 2025, *Configuring the PTP setting on the MOXA network switch*, Camera Technical Note CTN-003, NSF-DOE Vera C. Rubin Observatory, URL <https://ctn-003.lsst.io/>
- [SITCOMTN-136]**, Walter, C., 2024, *Calibration of the Rubin absolute timing*, Commissioning Technical Note SITCOMTN-136, NSF-DOE Vera C. Rubin Observatory, URL <https://sitcomtn-136.lsst.io/>

9 Acronyms

Acronym	Description
AST	NSF Division of Astronomical Sciences
AURA	Association of Universities for Research in Astronomy
CCS	Camera Control System
CTN	Camera Technical Note
DE-AC02	Department of Energy contract number prefix
FES	Filter Exchange System
LSST	Legacy Survey of Space and Time (formerly Large Synoptic Survey Telescope)
LSST-DA	LSST Discovery Alliance
OCS	Observatory Control System
PTC	Photon Transfer Curve
SB	Surface Brightness
SLAC	SLAC National Accelerator Laboratory