

Scheduling the LSST: Commissioning the Rubin Feature Based Scheduler

R. Lynne Jones^a, Peter Yoachim^b, Eric H. Neilsen, Jr.^c, Tiago Ribeiro^a, Erik Dennihy^a, and Sean Patrick MacBride^d

^aNSF-DOE Vera C. Rubin Observatory / NSF NOIRLab, 950 N. Cherry Ave., Tucson, AZ 85719, USA

^bUniversity of Washington, Dept. of Astronomy, Box 351580, Seattle, WA 98195, USA

^cFermi National Accelerator Laboratory, P. O. Box 500, Batavia, IL 60510, USA

^dPhysik-Institut, University of Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland

ABSTRACT

Over the ten years of the Legacy Survey of Space and Time (LSST), Rubin Observatory will acquire on the order of two million science observations. From twilight to twilight, except during bad weather or maintenance, the Simonyi Telescope will be choosing a new target every 30 seconds, with every observation contributing towards the LSST science goals. Planning these observations is the work of the Rubin Feature Based Scheduler (FBS). The FBS has been generating simulated observations for over ten years, helping to improve the LSST survey strategy and the algorithms within the FBS, in a collaborative process between the project and the astronomical community under the guidance of the LSST Survey Cadence Optimization Committee (SCOC). The FBS has been in active use with the Rubin Auxiliary Telescope for several years building the framework for interactions between the FBS and summit system. The commissioning of LSSTCam in 2025 marked the first on-sky use of the FBS with the Simonyi Telescope.

The commissioning of LSSTCam put the full capability of FBS into action with commissioning configurations similar to the full survey configuration. This configuration included multiple modes of observing, which automatically adapt to telemetry about the observatory, night sky and weather conditions during the night, varying from targeted pre-planned Deep Drilling Field observations to wide-area observations designed to acquire pairs of visits every 30 minutes. The FBS also appropriately responded to operational restrictions such as lunar avoidance zones as well as respect configurable altitude or azimuth constraints. During commissioning, we tested these FBS configurations, adjusting the survey strategy, monitoring FBS performance and survey progress, as well as learning important lessons about the LSST scheduling, the characteristics of resulting observations, and interactions with data management processing.

1. INTRODUCTION

The Vera C. Rubin Observatory began on-sky commissioning for the Legacy Survey of Space and Time (LSST) during 2025 and continued into early operations through mid-2026. Once full operations begin, the LSST will run continuously, every night, for ten years, observing in the *ugrizy* bandpasses. The survey addresses a broad range of science goals—from constraining dark energy and dark matter to mapping the Local Volume and Milky Way, cataloging small bodies in the Solar System, and exploring the transient and variable sky.¹ Achieving these diverse goals with a single survey requires an observatory purpose-built for the task and a carefully optimized observing strategy. Central to Rubin Observatory’s ability to deliver a wide, deep, and fast survey are its 6.5 m effective primary aperture, a 3.2 Gpixel camera with a 9.6 deg² field of view, and a mount that slews between pointings in under 5 s. Refining the observing strategy has relied on an extended program of simulations and community engagement to quantify how strategy changes affect scientific outcomes.

Scheduling decisions for the LSST, both in simulations and in real time at the summit, are complex and continuous. Each on-sky pointing must account for prior pointing history, current observatory status and performance, and prevailing weather and sky conditions. Decisions must both protect the facility (e.g., avoid pointing too close to the Moon) and maximize the scientific return of every exposure in service of the survey’s broad goals.

The decision-making process must be flexible enough to accommodate new constraints or evolving goals, and be capable of executing a ten-year simulation in a short time period to permit survey strategy optimization based on quantitative metrics developed by the project and the astronomical community. To meet these requirements, we have developed an automated scheduler called the Feature Based Scheduler (FBS). The FBS evolved from an earlier scheduler developed for initial LSST survey simulations.² Its current form reflects a subsequent redesign informed by Operations Research methods, better supporting the breadth and complexity of LSST scheduling.³ The FBS is implemented in the `rubin_scheduler` software package.⁴

This paper describes the Feature-Based Scheduler implementation (Section 2), the process for generating and analyzing simulation outputs (Section 3), early on-sky use of the scheduler (Section 4), lessons learned (Section 5), and concluding remarks (Section 6).

2. FEATURE BASED SCHEDULER IMPLEMENTATION

The Feature Based Scheduler (FBS) provides flexibility to support multiple scheduling modes tailored to the use case: it can follow a predefined sequence of observations that must be obtained under specified conditions, or it can select the best next exposure from approximately 1,500 candidate pointings across the visible sky by weighing factors such as expected image depth, slew time, and desired footprint coverage. This flexibility is provided through a range of `Survey` classes configured within the FBS `CoreScheduler`. The `CoreScheduler` acts as the framework that prioritizes and organizes the `Survey` instances it holds and provides the API to interact with observatory systems (or simulations)—ingesting telemetry on observatory state, proposing desired observations, and recording updates on acquired exposures.

2.1 Survey Classes

Within the `CoreScheduler`, `Survey` classes are configured for different goals or tasks and can be arranged in a priority order. The primary `Survey` classes are described below.

ScriptedSurvey. The `ScriptedSurvey` allows the definition of a predefined list of desired observations, each with constraints on the allowed time window, sky brightness, elevation, hour angle, lunar distance, and solar altitude. When the current conditions satisfy the constraints for any observation in the list, the `ScriptedSurvey` will request it.

Reward-based surveys (`BaseMarkovSurvey`). For surveys that must select optimal pointings across large areas of sky, the FBS provides the `BaseMarkovSurvey` base class and its subclasses. These surveys contain `BasisFunction` objects, which calculate the reward value of potential observations given the current `Conditions`. A simple example of a `BasisFunction` is one that assigns higher reward to shorter slews: it combines telemetry about the current telescope position with expected slew times to every other point on the sky and returns a reward that is a function of that slew time—not the slew time directly, but a value derived from it. Combining multiple `BasisFunction` instances with different weights yields a total reward. For example, combining a slew-time basis function with one that compares the current expected image depth to the best achievable depth and another that compares the current footprint coverage to the goal coverage produces a combined reward that balances uniform coverage, fast cadence, and photon collection per observation. Examples of these basis functions at a given time are illustrated in Figure 1. Subclasses of this base survey are used for the bulk of LSST observations, enabling optimization of the best pointings within a survey while also allowing prioritization between surveys.

FieldSurvey. The `FieldSurvey` is intended for use with a single chosen pointing and can execute a predefined sequence of observations each time it is selected, while still using basis functions to prioritize one `FieldSurvey` against another. A simple example is a `BasisFunction` that enforces a minimum revisit gap after the last visit to a given field before the survey can be chosen again, combined with a slew-time basis function that encourages shorter slews. If multiple `FieldSurvey` instances are configured with these basis functions, the FBS cycles through the fields at an interval equal to or longer than the revisit gap, obtaining the predefined observation sequence at each field’s location and moving to the next nearest field after each sequence.




Figure 1. The primary basis functions to balance image depth, uniform coverage, and on-sky efficiency are the M5DiffBasisFuntion, the SlewBasisFunction, and the FootprintBasisFunction, shown in the first three panels respectively. The stripe across the FootprintBasisFunction corresponds with the Galactic Plane, where the number of desired observations in this bandpass is lower than in the adjacent areas. The red triangle indicates the location of zenith, above Rubin, and in the SlewBasisFunction, the red X marks the current pointing of the telescope. These, combined with additional basis functions to fine tune survey visit choices and including the standard masks from Figure 2, are shown in the final panel. The final panel also shows the next pointing that would be chosen by this survey, as a red circle.

Target of Opportunity (ToOScriptedSurvey) The ToOScriptedSurvey class uses multiple inheritance from the above Survey classes respond to target of opportunity (ToO) events. The current configuration includes strategies for following up mulitple gravitational wave event triggers, neutrino detector triggers, solar system objects, and Galactic supernova. Most of the ToO follow up strategies were initially developed in Andreoni et al, 2022.⁵

2.2 Masks and Constraints

All Survey classes can accept BasisFunction instances that act as masks, removing inaccessible areas of sky or restricting available observation choices as necessary. This mechanism ensures that all surveys obey the same constraints and that the observatory operates safely within limits. The standard masks applied to all surveys include:

- Altitude and azimuth limits, including a zenith exclusion zone where tracking for the Simonyi Alt/Az mount is difficult.
- A Moon avoidance mask that ensures the telescope points no closer than 30 degrees from the Moon. Observing too close to the Moon risks tripping the LSSTCam HVBias, requiring over 30 minutes to reset. With more testing and experience regarding the exact behavior near this limit, we may be able to reduce the avoidance zone; for the moment, this keeps the camera safely operational.
- A planet avoidance mask that excludes Jupiter, Mars, and Venus, as these bright planets move quickly enough that it is preferable to avoid them rather than attempt to observe through them.

Figure 2 illustrates these standard masks at a specific time




Figure 2. Example sky map showing the standard masks applied to all **Survey** classes. Masked regions above the horizon are shown in brown and include altitude and azimuth limits (with a zenith exclusion zone), a 30 degree Moon avoidance zone, and exclusion regions around Jupiter and Venus (Mars is not visible). The unmasked area (light blue) represents the sky available for scheduling at this example epoch. The red triangle indicates the location of zenith, above Rubin.

During commissioning, an additional mask ensured that the Simonyi telescope remained pointed away from the rising Sun during the last three hours of the night, in case dome closure was delayed by system issues.

2.3 Operational Cycle

The FBS operates in a continuous cycle of three steps:

1. **Update telemetry** — `CoreScheduler.update_conditions()` ingests the current observatory and environmental state.
2. **Request observation** — `CoreScheduler.request_observation()` evaluates all configured surveys and returns the highest-priority feasible observation.
3. **Record observation** — `CoreScheduler.add_observation()` feeds the completed observation back into the scheduler, updating internal state.

This cycle repeats continuously while the FBS is in operation. Figure 3 illustrates this process.

2.4 Scheduler Configuration

A specific FBS configuration is defined by writing Python code that constructs a set of **Survey** instances and arranges them into prioritized tiers within the **CoreScheduler**. A higher-tier survey always takes priority over a lower-tier survey: lower tiers execute only when no survey in any higher tier has a feasible observation to request. Within each tier, surveys are ranked by their current reward values. Each **Survey** instance contains its own set of **BasisFunction** objects, which collectively define the constraints and optimization criteria that guide that survey's observation selection.

The tier structure and the number of surveys per tier are entirely determined by the use case. A simple configuration for a commissioning task or a targeted observing program might use a single tier with only one or two surveys. By contrast, a configuration for the full LSST survey contains on the order of ten tiers. This complexity supports a wide range of observing modes, including:




Figure 3. When the FBS is active, it continually cycles through the process of ingesting telemetry about the observatory conditions, returning requests for observations, and then ingesting information about acquired observations. When choosing observations in response to a request, the CoreScheduler iterates through the configured surveys. These surveys are arranged into tiers; surveys only compete within their own tier. If any survey within a given tier finds feasible observations, those observations will be chosen over any surveys in lower-ranking tiers. Within a given tier, the survey with the highest reward will request observations.

- **Target of Opportunity (ToO) observations**, in which a pre-defined sequence of observations must be executed if conditions permit in response to an external alert trigger.
- **Deep Drilling Field (DDF) observations**, in which a predefined sequence of visits are executed with a regular cadence, independently observing each of five DDF fields at appropriate times, ideally at the optimal time of night and during the desired phases of the lunar cycle, with the ability to catch up if observations are missed due to downtime or weather.
- **Wide-area observations acquired in pairs of visits**, where the two visits are separated by a predefined interval and acquired in a pre-specified combination of bandpasses. Pointing and bandpass selection must simultaneously balance depth, coverage, and efficiency (as described for reward-based surveys above) while distributing visits according to the desired coverage ratios between bandpasses and across the sky. There can be multiple tiers or surveys, where wide-area observations are configured with different requirements or different pair separations, or even without pairs at all.
- **Microsurveys with specialized requirements**, such as twilight observations of the ecliptic toward the Sun, which acquire a quad of visits per pointing rather than a pair.

This flexibility allows the same scheduling framework to serve purposes ranging from short on-sky commissioning sequences to the full LSST survey, and from real-time observing to ten year simulations.

3. SURVEY SIMULATIONS

The algorithms in the FBS are designed to make pointing decisions throughout the night under realistic observing conditions, while working towards scientific goals that can only be realized after much longer timeframes. To tune these algorithms and select the overall survey strategy that best satisfies the wide range of LSST science goals at the end of the ten year survey, the FBS is often run in a simulation mode using the same `rubin_scheduler` code base as in real-time operations, but with simulated telemetry replacing measurements from the summit and a model of the observatory’s response replacing actually acquired observations. This allows us to test variations on the scheduler configuration as well as generate predictions for survey performance into the future.

The FBS has been tested and developed against simulated telemetry and observation models since 2018, beginning with the LSST Call for White Papers on cadence optimization.⁶ Over 1000 simulated surveys have been produced to date, refining the scheduler configuration to balance science performance under the guidance of the Survey Cadence Optimization Committee (SCOC). The process of optimizing the survey strategy is documented in Bianco et al.⁷ and has involved significant input from the broader astronomical community. Quantitative analysis and comparison of these simulations is carried out with the `rubin_sim`⁸ MAF package.⁹ Simulations will continue throughout operations, responding to evolution in the scientific landscape and providing a means to evaluate the impact of observatory performance updates.

3.1 Sky and Atmospheric Models

To support survey strategy simulations, we have developed high-fidelity models of the expected sky conditions. An atmospheric seeing database based on historical DIMM measurements from CTIO¹⁰ is combined with the expected system contribution to seeing to produce wavelength- and airmass-dependent estimates of delivered image quality. A model of sky brightness as a function of time, bandpass, and position on the sky has been validated against on-site measurements at the summit.¹¹ Combining these models with the expected optical throughput of the telescope yields an estimate of the five-sigma point-source depth for each image.¹² Weather downtime is extrapolated from historical cloud-coverage records collected on-site.[?]

3.2 Observatory Performance Models

The model of observatory performance includes a high-fidelity model of the telescope itself, used to evaluate the slew time required between pointings. We have validated predictions from the slew model against measurements of the time required to move the Simonyi Telescope Mount Assembly (TMA) and found good agreement. We have also found that the total time required to complete slews during commissioning and early operations has




Figure 4. The slew model could accurately predict the time required to move the Simonyi Telescope Mount Assembly in altitude and azimuth, as shown by this comparison from a period during the SV survey, 2025-07-01 to 2025-07-10. However, the slew model was still somewhat inaccurate in predicting the entire visit gap. The offset of 1.6 seconds shown by the dotted red line can be explained by time required to configure and run scripts to execute the next step in the slew, but the scatter from that offset out to about 13.5 seconds is less clear. Similar tests run during the preLSST early operations phase show similar behavior.

involved additional overheads for telescope control processes and scripts to complete, adding some scatter in visit gaps that is not represented by our slew model. See Figure 4. Estimates of scheduled downtime due to planned engineering and maintenance, as well as unplanned downtime due to faults or technical issues, are also incorporated.

4. EARLY ON-SKY USE

At the summit, the FBS operates within the observatory’s Scheduler Commandable SAL Component (SchedulerCSC). The SchedulerCSC manages all interfaces with the observatory itself: gathering telemetry, commanding pointing requests, and handling the timing and queuing of observations selected by the FBS.

4.1 Auxiliary Telescope

Initial on-sky use of the FBS began with the Auxiliary Telescope (AuxTel)¹³ in early 2022. This early deployment exercised the end-to-end integration of the FBS into the SchedulerCSC: gathering telemetry, passing it to the FBS, receiving requested pointings, commanding the telescope, and reporting acquired observations back to the FBS. It also provided preliminary testing of the methods used to track how the FBS was configured and run and what artifacts were preserved to support identification of unanticipated scheduling choices.

However, AuxTel operation differed from the eventual LSST use case in several important respects. AuxTel’s data-acquisition rate was much lower and its instrumentation more varied (it was frequently operated with a spectrograph rather than an imager alone) so the cycle of requesting and reporting observations was considerably simpler. The slower cadence and longer exposures also gave observing specialists more opportunity and need to modify observations as they were acquired (e.g., changing the order or repeating a failed observation) outside of the FBS interface, and these ad hoc variations were difficult to capture in the recorded artifacts. The scheduling decisions themselves were also simpler: the target list was short and the selection algorithms bore little resemblance to those required for the full LSST survey. AuxTel use was therefore extremely valuable for integration testing but was not expected to be sufficient to fully commission the FBS for LSST survey-style observations.

4.2 Simonyi Telescope Commissioning

At the start of on-sky commissioning with LSSTCam on the Simonyi telescope, the FBS was used to execute simple observing programs consisting of a small number of fields observed through a fixed sequence of band-passes. Similar programs had previously been carried out with LSSTComCam.^{14,15} These tests served as further integration exercises for the FBS and SchedulerCSC on the Simonyi telescope, but the programs themselves were

deliberately simple: telemetry requirements were minimal, and observations for a given field were effectively triggered by the observing specialist pointing the telescope to the chosen field and then enabling the corresponding FBS configuration. This configuration, like the earlier AuxTel configurations, included only a collection of `FieldSurvey` instances, configured to acquire sequences of individual fields with the choice of field driven primarily by the Observing Specialists by including only a `SlewBasisFunction`.

The full Science Validation (SV) survey with LSSTCam began at the end of April 2025, marking the first deployment of an FBS configuration comparable in complexity to the full LSST survey. The SV survey plan is described in Clavar et al 2025¹⁶; in brief, approximately 3,000 deg² along the ecliptic were targeted for observations in a mode roughly approximating standard LSST Wide-Fast-Deep operations, albeit with a denser cadence and a smaller overall footprint. Planning for the SV survey was itself supported by simulations run with the FBS, which were used to position the 3,000 deg² footprint on the sky for optimal visibility during the SV observing period.

The FBS configuration used for the SV survey exercised most of the key features required for the full LSST survey. The configuration included multiple tiers, with reward-based surveys driving the bulk of the wide-area observations and `ScriptedSurvey` instances handling Deep Drilling Fields (DDFs). Operating in this mode provided commissioning experience that was not possible with the simpler programs used earlier. The FBS was now making genuine scheduling decisions, selecting the next pointing from a relatively large pool of candidates based on current conditions and survey history rather than executing a short, predetermined sequence on a single field. The success of the FBS and SchedulerCSC during the SV survey demonstrated their capability of scheduling the LSST during full operations.

4.2.1 Debugging with Artifacts

Debugging situations where the FBS did not behave as expected is supported by artifacts saved by the SchedulerCSC during on-sky operations at the summit. These artifacts include a snapshot of the telemetry (`Conditions`) passed to the FBS and a serialized copy of the FBS itself after ingesting that telemetry. Together, these snapshots allow us to recreate the FBS at the exact moment observations were requested and inspect its full internal state to understand why particular observations were chosen. The Python configurations for the FBS are maintained in a git repository deployed at the summit. The git hash of this repository, the version of the companion configuration-management package `ts_fbs_utils`, and the version of `rubin_scheduler` are all recorded in the LSST Engineering Facilities Database (EFD), providing a complete record of the scheduler configuration in use at any given time. Additionally, the SchedulerCSC’s list of upcoming observations (“targets”) and completed observations (“observations”) are also recorded in the EFD and can be compared against records of acquired observations generated from the images themselves. A few examples of issues found and resolved using these artifacts during this phase include:

- **Unexpected slew to an unrelated field.** After an initial sequence of visits, the FBS unexpectedly slewed away from the intended field. The targets and observations recorded in the EFD were reasonable and matched the on-sky behavior, but this was not the desired behavior (staying in the same location and taking a second sequence of visits of the intended field). It proved difficult to investigate the behavior from the snapshots and summit conditions available, as there was no snapshot for the second series of visits – only every snapshot was being saved. The survey scheduling team happened to be online and available during the night and began to investigate, but were hampered by the missing information. After several hours of work to attempt to fill in the gaps, we were able to recreate the unexpected behavior. The root cause was a configuration error: the field selection was driven by a slew-time basis function that was configured per-bandpass, and the final bandpass in the observation sequence had been omitted from that configuration by oversight. The FBS configuration was updated accordingly during the daytime, and a longer-term improvement to the underlying FBS basis functions to make configuration easier in the future was put in place with the next update of `rubin_scheduler` at the summit, several weeks later. In addition, the SchedulerCSC configuration was updated to record every snapshot, not just every other snapshot.
- **Incomplete bandpass sequence.** The FBS requested observations in several bandpasses, but only two were acquired. The FBS snapshot confirmed that the full sequence had been requested, yet the

recorded targets contained only a subset. Additional log messages recorded from the SchedulerCSC showed that the observations in the missing bandpass had been rejected due to filter change constraints. When SchedulerCSC experts were available, the issue was traced to a parameter in the SchedulerCSC that limited the maximum rate of filter changes for engineering reasons; this limit needed to be increased to permit the intended program to execute. This change was made and applied within the same night.

- **Dropped observations in a dither pattern.** A requested dither sequence was only partially acquired. The FBS snapshot again confirmed that the complete set of observations had been requested, but some did not appear in the recorded targets. Correlated records in the EFD revealed that observing-script failures immediately prior to the missing observations had caused them to be dropped, requiring adjustments in the SchedulerCSC’s handling of script failures. This was diagnosed during the next day after observing, and the fixes added at the next update of the SchedulerCSC.

Investigating these issues also led to improvements in the artifact-recording process itself. In particular, the cadence of FBS snapshots was increased from every other observation request to every request. The lower cadence had been adequate for AuxTel’s single-observation-at-a-time mode but was insufficient for the longer observation sequences used on the Simonyi telescope. Overall, the artifacts recorded by the FBS and SchedulerCSC have proven invaluable and have been sufficient to diagnose the majority of issues encountered to date.

It is worth noting that all of the issues described above arose from the relatively simple `FieldSurvey`-based configurations used during the earliest stages of commissioning. No comparable issues were observed during execution of the SV program. One Target of Opportunity (ToO) sequence failed to execute when desired, but this was quickly traced to the fact that the ToO target was closer than 30 degrees from the Moon: the FBS was behaving exactly as intended, keeping the observatory within predefined operational guidelines.

5. LESSONS LEARNED

The use of the FBS for the SV survey, as well as its subsequent use during Rubin’s early operations phase, has generally gone very smoothly. There are, of course, areas with room for improvement, either in the algorithms of the FBS itself or in its interactions with the SchedulerCSC. We have also gained additional first-hand knowledge of the precise needs for survey progress reporting and for downstream data processing.

5.1 Target of Opportunity Handling

One particular area where we identified room for improvement was in Target of Opportunity (ToO) coverage.

Using Rubin for ToO observations of well-localized sources is generally a poor use of resources, as these could be observed with other instruments with smaller fields of view. However, there may be cases where we do still receive a trigger for a ToO where the localization could fit within a single LSSTCam field of view, and efficient observing of such a ToO is important. During commissioning we triggered follow-up of the newly discovered interstellar comet 3I/ATLAS.¹⁷ We found that the framework for receiving alert locations in the FBS was not optimal for such well-localized events: the FBS attempted to cover not only the object’s position but also the surrounding `nside=32` HEALPixels, leading to inefficient coverage of the actual target. We have since added logic to the FBS and to the messages that pass ToO information into it, allowing well-localized objects to be identified as such so that the FBS can focus on the exact position.

Later, during the pre-LSST observing period, a ToO was triggered on gravitational-wave candidate S251112cm. This ToO covered an extremely large area of interest—larger than had been anticipated when designing the FBS ToO `ScriptedSurvey` templates. We found that the resulting coverage was inefficient: portions of the area of interest set below the horizon before the scheduler returned for subsequent cycles in different bandpasses. We have since added logic so that the FBS targets large ToO regions in smaller chunks ordered by Right Ascension, increasing the likelihood that the full area is covered before any portion sets.

5.2 Feedback from Data Management

The SV survey was the first opportunity to understand how observations acquired over a large area, using configurations similar to those planned for LSST operations, would interact with Rubin’s Data Management (DM) pipelines.

Earlier commissioning surveys had observed the DDF pointings, but with sequences unique to early commissioning. The SV survey was the first to observe the DDFs using the planned LSST sequences, which included small night-to-night dithers but no dithers within a single night for better efficiency in observing. We found that, without a large set of background images to smooth over the chip gaps, the lack of intra-night dithering produced image sets that DM could not effectively combine into coadds. We therefore added intra-night dithering alongside the existing night-to-night dithers at the cost of a short slew and settle period after every visit, both in the SV configuration and in the baseline LSST configuration.

The wide-field portion of the SV survey was also the first opportunity to test template generation using images obtained from the standard LSST sky tessellation and coverage pattern. We found that, due to the effects of chip gaps and the need for complete coverage over detector-sized patches on the sky, DM required more input images than anticipated in order to build template coadds. The acceptable images were also subject to tighter constraints on image quality (PSF and cloud extinction) than we had originally expected. Although much of this came to light only after the SV survey was well underway, the lesson has been folded back into the LSST survey strategy: we have added a “template tier” to the LSST configuration that prioritizes obtaining the minimum required number of images in regions of sky that do not yet meet the template goal, activated when atmospheric seeing is good. This increases the likelihood that the full sky will reach the required number of high-quality images in each bandpass.

5.3 Comparison of Simulations to On-Sky Performance

The SV survey provided a valuable opportunity to compare simulated observations against those actually acquired. We were able to validate our model telemetry inputs, and even use the models to identify performance issues at the summit. By comparing the predicted visit gaps based on the slew model and expected overheads to the actual visit gaps, we could identify periods where the expected overheads had increased, such as when the SchedulerCSC was only configuring a single script at a time instead of configuring the next two scripts ahead of time.

Comparing more detailed behavior of the FBS from simulations to reality proved challenging. The principal difficulty was that the SV survey was not the sole activity on the telescope: large portions of many nights were spent on engineering tests, and additional nights were lost to weather, so the periods of true SV observing were neither continuous nor predictable in advance.

Using a retrospective approach, running simulations only over the intervals during which the SV program was demonstrably active on-sky, we were able to reproduce the overall number of visits to within roughly 10%, with a consistent general pattern of coverage on the sky. The remaining discrepancy can be attributed to several factors: scatter in the visit gap (as discussed in Section 3.2); short faults that fell below our threshold for ending a continuous observing interval but nevertheless reduced on-sky efficiency; and differences in fault recovery between the simulation and real operations, in particular the simulation’s assumption that the telescope returns to a parked position after a fault rather than remaining at the last pointing. This is an area we continue working on improvement. Additional telemetry from the summit about the observatory state, including weather closures or periods of fault, is now available and will provide useful inputs.

5.4 Survey Progress Reporting

A recurring theme throughout the SV survey and early operations has been the need to communicate survey progress and near-term plans to the broader Rubin science community. We have iterated on several reporting mechanisms, each of which has informed our understanding of what the community finds most useful.

Nightly reports summarizing the visits acquired during the previous night were generated automatically from rendered Jupyter notebooks and published shortly after each night’s observing. These reports provided a consistent, low-latency view of survey progress and were well received. When specific feedback did reach us, we

were able to act on it. For example, a community member asked whether we could report the number of visits acquired in pairs each night, and we added this to the nightly reports.

In parallel, we published less frequent but more in-depth reports on a dedicated website. These were authored by hand and could move beyond the bare record of acquired visits to discuss periods of bad weather, changes in commissioning goals, and other context that does not fit naturally into an automated notebook. These reports also included FBS simulations projecting likely coverage for the remainder of the SV period, giving readers a sense of where the survey was heading rather than only where it had been. Like the nightly reports, these in-depth summaries were very well received.

Looking ahead, we plan to add a third class of report describing the upcoming night: the planned configuration, the regions of sky likely to be observed, and the conditions under which various contingencies would apply. We expect this forward-looking view to be particularly useful for community members planning coordinated observations or follow-up programs, and for setting realistic expectations about which fields will be revisited on a given timescale.

Establishing more effective channels for community feedback on all three classes of report—nightly, in-depth, and forward-looking—remains an open area for improvement. Keeping the reports useful to the broad Rubin community requires balancing the specific needs of individual users against the readability of the report as a whole—a balance that is easier to strike with a structured feedback process than with ad hoc suggestions.

6. CONCLUSION

The Feature-Based Scheduler has matured from a simulation tool used to explore LSST cadence strategies into the operational scheduler driving the Vera C. Rubin Observatory on the sky. Its design—a configurable set of **Survey** classes coordinated by the **CoreScheduler**, with reward calculations built from composable **BasisFunction** objects—has proven flexible enough to support a wide range of observing modes, from short commissioning sequences on a single field to the full complexity of the Science Validation survey and, ultimately, the ten-year LSST.

Years of simulation work, carried out in close collaboration with the Survey Cadence Optimization Committee and the broader astronomical community, allowed the FBS to be tuned against realistic models of the sky, atmosphere, and observatory long before first light. On-sky commissioning, first with the Auxiliary Telescope and then with the Simonyi telescope and LSSTCam, validated the integration of the FBS with the SchedulerCSC and exercised progressively more of the scheduler’s capabilities, culminating in the SV survey’s demonstration of LSST-style operations.

Early operational experience has also identified concrete areas for continued development: improved handling of well-localized and very large Target of Opportunity regions, configuration changes informed by feedback from the Data Management pipelines (intra-night DDF dithers and the addition of a template tier), and ongoing work to communicate survey progress and plans to the community in a way that is both informative and sustainable. These refinements will continue throughout operations as the scientific landscape evolves and as observatory performance is better characterized on the sky.

As part of the commissioning surveys that ran from May to September 2025, the FBS drove the acquisition of over 20,000 science visits, including over 13,000 designed to be similar to the Wide Fast Deep survey, just under 1000 visits designed to be like Deep Drilling Fields, and almost 200 visits to test ToO acquisition. In the period of early operations from October 2025 to present, the FBS has been responsible for choosing almost 50,000 additional visits. With the FBS now scheduling LSST observations in real time, the framework described in this paper provides the foundation for ten years of survey operations and for the continued optimization of Rubin’s scientific return.

ACKNOWLEDGMENTS

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.

REFERENCES

- [1] Ivezić, Ž., Kahn, S. M., Tyson, J. A., Abel, B., Acosta, E., Allsman, R., Alonso, D., AlSayyad, Y., Anderson, S. F., Andrew, J., Angel, J. R. P., Angeli, G. Z., Ansari, R., Antilogus, P., Araujo, C., Armstrong, R., Arndt, K. T., Astier, P., Aubourg, É., Auza, N., Axelrod, T. S., Bard, D. J., Barr, J. D., Barrau, A., Bartlett, J. G., Bauer, A. E., Bauman, B. J., Baumont, S., Bechtol, E., Bechtol, K., Becker, A. C., Becla, J., Beldica, C., Bellavia, S., Bianco, F. B., Biswas, R., Blanc, G., Blazek, J., Blandford, R. D., Bloom, J. S., Bogart, J., Bond, T. W., Booth, M. T., Borgland, A. W., Borne, K., Bosch, J. F., Boutigny, D., Brackett, C. A., Bradshaw, A., Brandt, W. N., Brown, M. E., Bullock, J. S., Burchat, P., Burke, D. L., Cagnoli, G., Calabrese, D., Callahan, S., Callen, A. L., Carlin, J. L., Carlson, E. L., Chandrasekharan, S., Charles-Emerson, G., Chesley, S., Cheu, E. C., Chiang, H.-F., Chiang, J., Chirino, C., Chow, D., Ciardi, D. R., Claver, C. F., Cohen-Tanugi, J., Cockrum, J. J., Coles, R., Connolly, A. J., Cook, K. H., Cooray, A., Covey, K. R., Cribbs, C., Cui, W., Cutri, R., Daly, P. N., Daniel, S. F., Daruich, F., Daubard, G., Daues, G., Dawson, W., Delgado, F., Dellapenna, A., de Peyster, R., de Val-Borro, M., Digel, S. W., Doherty, P., Dubois, R., Dubois-Felsmann, G. P., Durech, J., Economou, F., Eifler, T., Eracleous, M., Emmons, B. L., Fausti Neto, A., Ferguson, H., Figueroa, E., Fisher-Levine, M., Focke, W., Foss, M. D., Frank, J., Freemon, M. D., Gangler, E., Gawiser, E., Geary, J. C., Gee, P., Geha, M., Gessner, C. J. B., Gibson, R. R., Gilmore, D. K., Glanzman, T., Glick, W., Goldina, T., Goldstein, D. A., Goodenow, I., Graham, M. L., Gressler, W. J., Gris, P., Guy, L. P., Guyonnet, A., Haller, G., Harris, R., Hascall, P. A., Haupt, J., Hernandez, F., Herrmann, S., Hileman, E., Hoblitt, J., Hodgson, J. A., Hogan, C., Howard, J. D., Huang, D., Huffer, M. E., Ingraham, P., Innes, W. R., Jacoby, S. H., Jain, B., Jammes, F., Jee, M. J., Jenness, T., Jernigan, G., Jevremović, D., Johns, K., Johnson, A. S., Johnson, M. W. G., Jones, R. L., Juramy-Gilles, C., Jurić, M., Kalirai, J. S., Kallivayalil, N. J., Kalmbach, B., Kantor, J. P., Karst, P., Kasliwal, M. M., Kelly, H., Kessler, R., Kinnison, V., Kirkby, D., Knox, L., Kotov, I. V., Krabbendam, V. L., Krughoff, K. S., Kubánek, P., Kuczewski, J., Kulkarni, S., Ku, J., Kurita, N. R., Lage, C. S., Lambert, R., Lange, T., Langton, J. B., Le Guillou, L., Levine, D., Liang, M., Lim, K.-T., Lintott, C. J., Long, K. E., Lopez, M., Lotz, P. J., Lupton, R. H., Lust, N. B., MacArthur, L. A., Mahabal, A., Mandelbaum, R., Markiewicz, T. W., Marsh, D. S., Marshall, P. J., Marshall, S., May, M., McKercher, R., McQueen, M., Meyers, J., Migliore, M., Miller, M., and Mills, D. J., “LSST: From Science Drivers to Reference Design and Anticipated Data Products,” *ApJ* **873**, 111 (March 2019). DOI: <https://doi.org/10.3847/1538-4357/ab042c>.
- [2] Delgado, F. and Reuter, M. A., “The LSST Scheduler from design to construction,” in [*Observatory Operations: Strategies, Processes, and Systems VI*], Peck, A. B., Seaman, R. L., and Benn, C. R., eds., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **9910**, 991013 (July 2016). DOI: <https://doi.org/10.1117/12.2233630>.
- [3] Naghib, E., Yoachim, P., Vanderbei, R. J., Connolly, A. J., and Jones, R. L., “A Framework for Telescope Schedulers: With Applications to the Large Synoptic Survey Telescope,” *AJ* **157**, 151 (Apr. 2019). DOI: <https://doi.org/10.3847/1538-3881/aafecf>.
- [4] Yoachim, P., Jones, L., Eric H. Neilsen, J., MacBride, S., Bechtol, K., Becker, M. R., Ross, Espada, G., and humna, “lsst/rubin_scheduler: v3.22.0,” (Jan. 2026). <https://doi.org/10.5281/zenodo.18395761>.
- [5] Andreoni, I., Margutti, R., Salafia, O. S., Parazin, B., Villar, V. A., Coughlin, M. W., Yoachim, P., Mortensen, K., Brethauer, D., Smartt, S. J., Kasliwal, M. M., Alexander, K. D., Anand, S., Berger, E., Bernardini, M. G., Bianco, F. B., Blanchard, P. K., Bloom, J. S., Brocato, E., Bulla, M., Cartier, R., Cenko, S. B., Chornock, R., Copperwheat, C. M., Corsi, A., D’Ammando, F., D’Avanzo, P., Hélène Datrier, L. É., Foley, R. J., Ghirlanda, G., Goobar, A., Grindlay, J., Hajela, A., Holz, D. E., Karambelkar, V., Kool, E. C., Lamb, G. P., Laskar, T., Levan, A., Maguire, K., May, M., Melandri, A., Milisavljevic, D., Miller, A. A., Nicholl, M., Nissanke, S. M., Palmese, A., Piranomonte, S., Rest, A., Sagués-Carracedo, A., Siellez, K., Singer, L. P., Smith, M., Steeghs, D., and Tanvir, N., “Target-of-opportunity Observations of Gravitational-wave Events with Vera C. Rubin Observatory,” *ApJS* **260**, 18 (May 2022). DOI: <https://doi.org/10.3847/1538-4365/ac617c>.

- [6] Ivezić, Ž., Jones, L., Ribeiro, T., LSST Project Scientist Team, and LSST Science Advisory Committee, “Call for White Papers on LSST Cadence Optimization,” Informal Construction Document Document-28382, NSF-DOE Vera C. Rubin Observatory (June 2018). <https://ls.st/Document-28382>.
- [7] Bianco, F. B., Ivezić, Ž., Jones, R. L., Graham, M. L., Marshall, P., Saha, A., Strauss, M. A., Yoachim, P., Ribeiro, T., Anguita, T., Bauer, A. E., Bauer, F. E., Bellm, E. C., Blum, R. D., Brandt, W. N., Brough, S., Catelan, M., Clarkson, W. I., Connolly, A. J., Gawiser, E., Gizis, J. E., Hložek, R., Kaviraj, S., Liu, C. T., Lochner, M., Mahabal, A. A., Mandelbaum, R., McGehee, P., Neilsen, Jr., E. H., Olsen, K. A. G., Peiris, H. V., Rhodes, J., Richards, G. T., Ridgway, S., Schwamb, M. E., Scolnic, D., Shemmer, O., Slater, C. T., Slosar, A., Smartt, S. J., Strader, J., Street, R., Trilling, D. E., Verma, A., Vivas, A. K., Wechsler, R. H., and Willman, B., “Optimization of the Observing Cadence for the Rubin Observatory Legacy Survey of Space and Time: A Pioneering Process of Community-focused Experimental Design,” *ApJS* **258**, 1 (January 2022). DOI: <https://doi.org/10.3847/1538-4365/ac3e72>.
- [8] Yoachim, P., Jones, L., Eric H. Neilsen, J., Tiago, Leistedt, B., Parejko, J., Bellm, E., Street, R., Carlin, J., humna, Becker, M. R., pgris, erykoff, reneehlozek, Prinszano, L., Gollotti, G. A., Dennihy, E., Sick, J., lmptc, Abrams, N., Ross, Assef, R. J., LI, Bregeon, J., Lim, K.-T., Guy, L., Kelley, M., and Andreoni, I., “lsst/rubin_sim: v2.6.2,” (May 2026). <https://doi.org/10.5281/zenodo.20451985>.
- [9] Jones, R. L., Yoachim, P., Chandrasekharan, S., Connolly, A. J., Cook, K. H., Ivezić, Ž., Krughoff, K. S., Petry, C., and Ridgway, S. T., “The LSST metrics analysis framework (MAF),” in [*Observatory Operations: Strategies, Processes, and Systems V*], Peck, A. B., Benn, C. R., and Seaman, R. L., eds., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **9149**, 91490B (July 2014). DOI: <https://doi.org/10.1117/12.2056835>.
- [10] Neilsen, E., “Seeing values for LSST strategy simulations,” Technical Note RTN-022, NSF-DOE Vera C. Rubin Observatory (July 2023). <https://rtn-022.lsst.io/>.
- [11] Yoachim, P., Coughlin, M., Angeli, G. Z., Claver, C. F., Connolly, A. J., Cook, K., Daniel, S., Ivezić, Ž., Jones, R. L., Petry, C., Reuter, M., Stubbs, C., and Xin, B., “An optical to IR sky brightness model for the LSST,” in [*Observatory Operations: Strategies, Processes, and Systems VI*], Peck, A. B., Seaman, R. L., and Benn, C. R., eds., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **9910**, 99101A (July 2016). DOI: <https://doi.org/10.1117/12.2232947>.
- [12] Jones, R. L., “Calculating LSST limiting magnitudes and SNR,” Simulations Team Technical Note SMTN-002, NSF-DOE Vera C. Rubin Observatory (July 2025). <https://smtn-002.lsst.io/>.
- [13] Ingraham, P., Clements, A. W., Ribeiro, T., Reuter, M. A., Fisher-Levine, M., Hoblitt, J., Lupton, R. H., Thomas, S., Stubbs, C. W., Arndt, K. T., Callahan, N., Claver, C. F., Colleoni, F., Corral, L., Doherty, P., Economou, F., Fausti Neto, A., Jenness, T., Khine, H., Krughoff, K. S., Mondrik, N., Menanteau, F., Mills, D. J., O’Mullane, W., Reil, K. A., Rivera, M., Stalder, B., Sebag, J., Shipsey, I., Tighe, R., Thornton, A. J., Villalobos, A., and Wiecha, O., “Vera C. Rubin Observatory auxiliary telescope commissioning as a control system pathfinder,” in [*Software and Cyberinfrastructure for Astronomy VI*], Guzman, J. C. and Ibsen, J., eds., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **11452**, 114520U (Dec. 2020). DOI: <https://doi.org/10.1117/12.2561112>.
- [14] Vera C. Rubin Observatory, “An Interim Report on the LSSTComCam On-Sky Campaign,” Commissioning Technical Note SITCOMTN-149, NSF-DOE Vera C. Rubin Observatory (August 2025). <https://sitcomtn-149.lsst.io/>.
- [15] Stalder, B., Reil, K., Aguilar, C., Araya, C., Borstad, A., Bowdish, B., Cho, M., Cisneros, S., Claver, C., Constanzo, J., Corvetto, G., Drass, H., Economou, F., Dennihy, E., Eisner, A., Fabrega, J., Fausti Neto, A., Fisher-Levine, M., Gonzalez, I., Harris, R., Hascall, D., Haupt, J., Hoblitt, J., Howard, J., Huffer, M., Jenness, T., Jiménez Mejías, D., Johnson, B., Johnson, T., Krughoff, K. S., Lage, C., Lange, T., Liang, M., Lopez, J., Lopez, M., Lupton, R., Maulen, G., Menanteau, F., Mills, N., Morris, G. R., Munoz, F., Neal, H., Neill, D. R., Newbry, S., Nomerotski, A., Onoprienko, D., Ordenes, I., Orellana, J., Osier, S., Park, H., Pietrowicz, S., Plazas, A., Poczulp, G., Qiu, B., Quint, B., Reinking, H., Reuter, M., Ribeiro, T., Rojas, R., Romero, S., Schindler, R., Schoening, B., Sebag, J., Shugart, A., Silva, C., Sotuela, I., Tache, A., Tapia, D., Thayer, J. G., Thomas, S., Tighe, R., Tsai, T.-W., Turri, M., Tyson, T., Vergara, L., Walter, C., Wiecha, O., and Xiong, V., “Rubin Observatory Commissioning Camera: summit integration,” in [*Ground-based and Airborne Instrumentation for Astronomy IX*], Evans, C. J., Bryant, J. J., and Motohara, K., eds., *Society of*

Photo-Optical Instrumentation Engineers (SPIE) Conference Series **12184**, 121840J (August 2022). DOI: <https://doi.org/10.1117/12.2630184>.

- [16] Claver, C., Bauer, A., Bechtol, K., Bellm, E., Blum, R., Bosch, J., Clements, A., Connolly, A., Guy, L., Ivezić, Ž., Lupton, R., Ritz, S., O'Mullane, W., Thomas, S., Drass, H., and Jones, L., "Construction Completeness and Operations Readiness Criteria," Commissioning Technical Note SITCOMTN-005, NSF-DOE Vera C. Rubin Observatory (June 2025). <https://sitcomtn-005.lsst.io/>.
- [17] Chandler, C. O., Bernardinelli, P. H., Jurić, M., Singh, D., Hsieh, H. H., Sullivan, I., Jones, R. L., Kurlander, J. A., Vavilov, D., Eggl, S., Holman, M., Spoto, F., Schwamb, M. E., MacArthur, L. A., Makadia, R., Micheli, M., Heinze, A., Christensen, E. J., Beebe, W., Roodman, A., Lim, K.-T., Jenness, T., Bosch, J., Smart, B. M., Bellm, E., MacBride, S., Rawls, M. L., Greenstreet, S., Slater, C., Ivezić, Ž., Blum, R. D., Connolly, A., Daues, G., Gower, M., Kalmbach, J. B., Bannister, M. T., Dones, L., Dorsey, R. C., Farnocchia, D., Fraser, W. C., Forbes, J. C., Fuentes, C., Holt, C. E., Inno, L., Jones, G. H., Knight, M. M., Lintott, C. J., Lister, T., Lupton, R., Magbanua, M. J. M., Malhotra, R., Mueller, B. E. A., Murtagh, J., Pandey, N., Reach, W. T., Samarasinha, N. H., Seligman, D. Z., Snodgrass, C., Solontoi, M., Szabó, G. M., Vereš, P., White, E., Womack, M., Young, L. A., Allbery, R., Anand, S., Armellini, R., Aubourg, É., Avdellidou, C., Azfar, F., Bauer, J., Bechtol, K., Becker, V., Belyakov, M., Benecchi, S. D., Bertini, I., Bodewits, D., Boeshaar, P., Bolin, B. T., Bose, M., Boucaud, A., Bouffleur, R. C., Boutigny, D., Bradshaw, A., Braga-Ribas, F., Bregeon, J., Buchanan, L. E., Calabrese, D., Camargo, J. I. B., Caplar, N., Carlin, J. L., Carry, B., Carvajal, J. P., Ceballo, R., Chiang, H.-F., Choi, Y., Combet, C., da Costa, L., Cowan, P., Crenshaw, J. F., Croft, S., Čuk, M., Daly, P. N., D'Ammando, F., Daruich, F., Daubard, G., Davenport, J. R. A., Daylan, T., Delgado, J., Devillepoix, H. A. R., Doherty, P. E., Donaldson, A., Drass, H., Deppe, S. J. H., Dubois-Felsmann, G. P., Ferguson, P. S., Economou, F., Eduardo, M. R., Sotuela Elorriaga, I., Englert, A., Karavakis, E., Fanning, K., Frissell, M. K., Fedorets, G., Fernandes, M. B., Ferté, A., Fisher-Levine, M., Freytag, M. L., Fulle, M., Gandhi, P., Gates, J., Gerdes, D. W., Gibbs, A. R., Gillan, A. F., Giordano Orsini, M., Glanzman, T., Goodenow, I., Gomes-Júnior, A. R., Gorsuch, M. R., Granvik, M., Guan, W., Guy, L. P., Hammergren, M., Hanushevsky, A., Hernandez, F., Herrold, Å., Hestroffer, D., Hoblitt, J., Megias Homar, G., Hopkins, M. J., Ieva, S., Ingraham, P., Irving, D. H., Jannuzi, B. T., Jee, M. J., Jimenez, D., Juramy, C., Kahn, S. M., Kang, Y., Kannawadi, A., Kavelaars, J., Kelley, M. S. P., Kelkar, K., Kelvin, L. S., Kotov, I., Koumjian, A., Kovács, G., Krughoff, K. S., Kryszczyńska, A., Kubánek, P., Lage, C., Lange, T. J., Léget, P.-F., Le Guillou, L., Levine, B., Levine, W. G., Li, Z. C., Liang, S., Licandro, J., Lin, H. W., Lisse, C., Lust, N. B., Lyttle, R. R., Mahabal, A. A., Mahlke, M., Mainetti, G., Mandelbaum, R., Margheim, S. J., Margoti, G., Marshall, P., Plazas Malagón, A. A., Marčeta, D., Melita, M. D., Menanteau, F., Meyers, J., Mills, D., Moniez, M., Morales Marín, C. A. L., Morato, N., and More, S., "NSF-DOE Vera C. Rubin Observatory Observations of Interstellar Comet 3I/ATLAS (C/2025 N1)," *ApJL* **1001**, L35 (Apr. 2026). DOI: <https://doi.org/10.3847/2041-8213/ae4b3a>.