



Vera C. Rubin Observatory
Project Science Team

Survey Cadence Optimization Committee's Phase 3 Recommendations

the SCOC

PSTN-056

Latest Revision: 2024-09-07

DRAFT



Abstract

The final recommendation of the SCOC for the Survey Strategy prior to the start of LSST

Draft

Change Record

Version	Date	Description	Owner name
1	YYYY-MM-DD	Unreleased.	Federica Bianco and the Survey Cadence Optimization Committee

Document source location: <https://github.com/lstt-pst/pstn-056>

Draft

Contents

1 Introduction	1
2 Executive Summary of the Phase 3 recommendation	2
2.1 Note on how to read the SCOC plots	2
3 Open questions from PSTN-055	4
3.1 Swapping filters on the filter wheel	4
3.2 Filter Balance	6
3.3 Rolling	11
3.4 Galaxy	16
3.4.1 Footprint and Time Distribution of Visits	16
3.4.2 Galactic Filter Balance	19
3.4.3 The Milky Clouds and South Celestial Pole	20
3.5 Targets of Opportunity (ToO)	21
3.6 Snaps	24
3.7 Deep Drilling Fields (DDF)	25
3.8 Early Science	26
4 Additional Recommendations	27
5 Additional changes introduced in v3.6 OpSims	28
6 Summary of SCOC Phase 3 recommendations	28
A References	34
B Acronyms	35

Survey Cadence Optimization Committee's Phase 3 Recommendations

1 Introduction

With an unprecedented engagement of the scientific community at large, the Vera C. Rubin Observatory (hereafter Rubin) has designed a process of incremental improvements on the survey strategy to maximize the overall scientific throughput of the Legacy Survey of Space and Time (LSST). The high-level requirements for the LSST are set by four science pillars: probing dark energy and dark matter, building an unprecedented inventory of the Solar System, mapping the Milky Way and Local Volume, and exploring the transient universe. These requirements are described in Ivezić & The LSST Science Collaboration (LPM-17) (hereafter Science Requirements Document, or SRD, but significant flexibility in survey cadence remains within these requirements. The optimization of the survey strategy process is aimed at maximizing science for the four science pillars and increase the portfolio of LSST science by tuning the survey strategy and cadence within the SRD requirements (SRD).

As part of this process, the Survey Cadence Optimization Committee (SCOC) was set up by Rubin's Science Advisory Committee in 2018 to solicit review and integrate community feedback at large and make recommendations for the implementation of the LSST survey strategy to the Director of Operations. This document constitutes the third SCOC recommendation, resulting from the phase 3 process of survey design which started in January 2023, after the delivery of the Phase 2 recommendation (F. Bianco and The Rubin Observatory Survey Cadence Optimization Committee PSTN-055 —hereafter PSTN-055— and the baseline simulation baseline_v3.0. This will be the last recommendation leading to the start of LSST. Yet, reviews of the survey strategy will continue throughout the 10-years of LSST with the SCOC reviewing the survey throughput and community feedback and renewing its recommendation on an annual basis.

The Phase 3 recommendation responds directly to the questions left open in PSTN-055 and updates and refines previous recommendations (PSTN-055 and Z. Ivezić and The Rubin Observatory Survey Cadence Optimization Committee (PSTN-053), hereafter PSTN-053) The present document generally does not reiterate previous recommendations that have not changed.

Finally, to help the reader understand the text that follows, we note that the LSST survey is

actually an ensemble of surveys. It includes a main survey, known as Wide Fast Deep (WFD) which by SRD requirement should receive more than 825 observations and covered at least 18,000 square degrees, a Galactic Plane (and Bulge) survey. Furthermore, a WFD low-dust region is defined with limits $-70^\circ \leq \text{Dec} \leq +12.5^\circ$ for $RA \sim 7 - 18h$ and $-72^\circ \leq \text{Dec} \leq +3^\circ$ for $0 \lesssim RA \lesssim 7h$ and $18h \lesssim RA \lesssim 24h$, with the addition of the Virgo cluster PSTN-053. Special regions in the LSST include the LMC, SMC, and South Celestial Pole and North Ecliptic Spur. Microsurveys can be performed and PSTN-053 committed to perform a twilight Near-Earth-Objects microsurvey in Y1 and to accept further proposals for future microsurveys after the start of LSST.

The document is structured as follows.

The open questions identified in PSTN-055 Section 4 are included here for the reader's convenience in section 3. The SCOC recommendations relative to those questions follow each question.

Additional changes to the survey strategy are described in section 4.

The current recommendation is summarized in section 6.

2 Executive Summary of the Phase 3 recommendation

Notable updates from previous recommendations, and corresponding changes to the baseline, include: updating system throughputs, expectations for engineering time, and slew performance; tweaking the filter balance in response to throughput changes; refining observing choices in the Galactic Plane, Bulge, and LMC/SMC/SCP; defining the implementation plan for the ToO program; recommending single visit exposures over visits implemented in "snaps"; investigating new rolling strategy options; refining the DDF observing plans.

2.1 Note on how to read the SCOC plots

In this document you will see sky maps measuring quantities (e.g., number of visits) in healpixels. The typical sky pixelization that underlies the metric calculations the SCOC review is 128 sides healpixels (covering an area $\sim 0.2\text{deg}^2$) although for particularly computationally intense MAFs this is turned down to 64 or 32.

The SCOC typically reviews the outcome of MAFs across multiple OpSim to compare scientific performance. You will see plots in three styles:

Heatmaps: Figure 1 - the divergent color scheme shows improvements in metrics in blue and drops in metrics in red. Note that different heat maps may show different ranges in the color scheme, but the SCOC typically considers changes of more than a few percent to be significant and less than a few percent to be in the noise. One of the OpSim is chosen as comparison and the corresponding column will look neutral in color.




FIGURE 1: A heatmap produced by the Observing Strategy team for the SCOC. The plot compares the performance of selected metrics across baseline simulations from baseline_v1.X (original survey strategy) through baseline_v3.4. In these plots, a OpSim is selected as the reference and all other OpSims' performances are shown relative to that. That is: the reference OpSim (baseline_v2.0 in this case) has MAF=1 for all the metrics. Blue colors indicate a positive metric value, *i.e.* an improvement. Red colors indicate a performance drop with respect to the reference OpSim.

Radar plots: Figure 2 - when comparing small numbers of metrics and few OpSim we often use radar plots. Each corner of the radar plot corresponds to a metric and the colored lines inside of the plots that join each corner show metric performance. In these plots the comparison OpSims looks like a N-gone (or N-sided circle). Where the MAF performance shows improvements compared to the comparison OpSim the point lies outside of the N-gone, where its a loss it sits inside. The range of performance changes, so make sure you carefully inspect the

plot to see the performance scaling going in and out of the N-gone (typical values are 0.9-1.1).

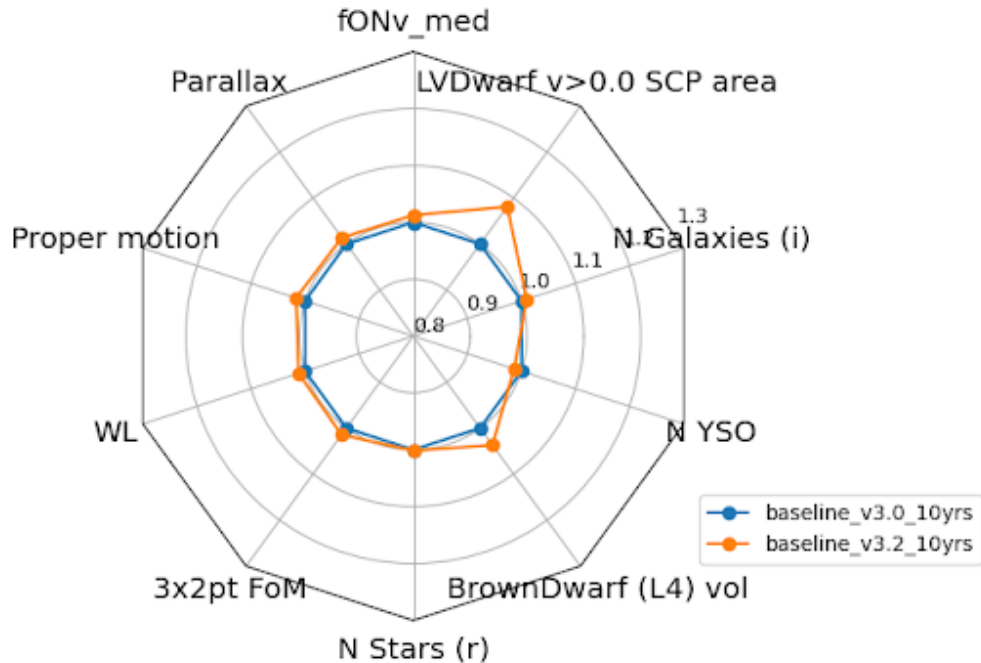


FIGURE 2: A radar plot comparing performance of MAFs under different filter swapping schemes (see subsection 3.1). This plot compares baseline_v3.0 with baseline_v3.2. All metrics within perform as well or better in baseline_v3.2 as shown by the orange point laying outside of the blue polygon, except N YSO where, however, the loss is minimal and likely not statistically significant. The range of the axis is 0.8-1.3, indicating that a point lying in the center would measure a performance 20% worse than the reference OpSim and a point on the outer perimeter would indicate a 30% improvement.

3 Open questions from PSTN-055

3.1 Swapping filters on the filter wheel

The filter system at Rubin allows five of the six filters to be mounted at the same time on the carousel. Filters at the edge of the system transmission will be swapped in and out of the filter wheel throughout the lunar cycle. In PSTN-055 Section 4 the SCOC recommended further investigation of which filters to swap:

[PSTN-055] “The SCOC recommends that the investigation of the filter swapping schemes on the filter wheel continue. After the November 2022 SCOC workshop, a few experiments in swapping u , z , and y instead of u and z were implemented in

v2.99 simulations. More work is needed to understand the impacts of this decision on the DDFs as well as on the WFD. Filter pairing prescriptions for the observation pairs should also be explored in some more depth.”

Simulations prior to v3.0 swapped the z with u based on lunation, as moonlight is blue and impacts observations in u -band most significantly. Simulations tagged v3.2 experimented with swapping u with z or y , including putting all of u , z , and y on rotation. Increasing the availability of z on the filter wheel produced significant improvements in supernova (SN) cosmology, especially in the Deep Drilling Fields (DDFs), while swapping two filters, instead of three, improves coverage at short time scales in filters through z (see Figure 3).

The SCOC recommends swapping u - and y -band according to the moon phase. This produces benefits for SN cosmology while preserving coverage on short timescales. This recommendation is implemented starting in baseline_v3.2.

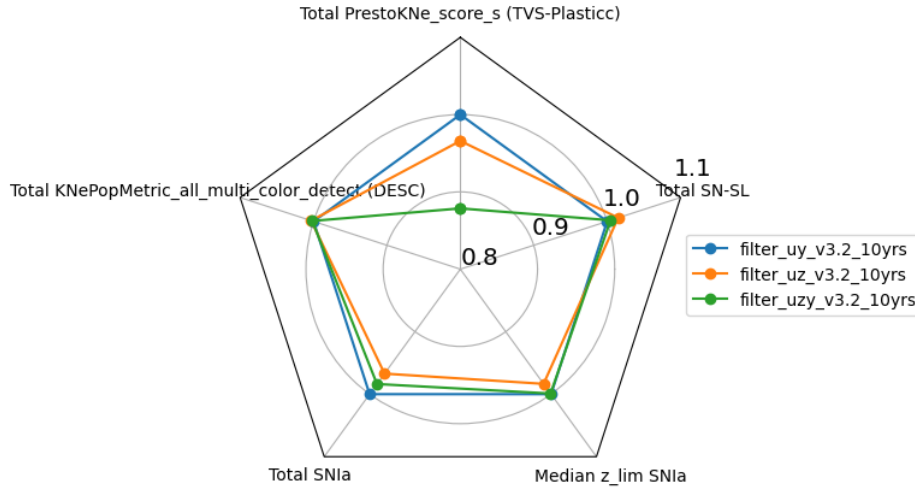


FIGURE 3: Swapping a filter out of the camera has two effects: it adds a gap for the period of time while it’s unavailable, and it increases the cadence in that bandpass during the time it is in the camera. Gaps in z band are particularly bad for SNIa detections, an effect which is magnified for Deep Drilling Fields. Swapping three filters (u with alternating z and y) increases the length of time gaps between sampling in the same filter, decreasing performance for hours-days transients. Swapping u with y while leaving z mounted on the filter wheel has an overall positive science impact, balancing the needs of rapid transient science with SNIa science in the Deep Drilling Fields.

3.2 Filter Balance

In PSTN-055, the SCOC confirmed the recommendation on the filter balance as implemented starting in `baseline_v2.0` but left the possibility that:

[PSTN-055] “While the SCOC recommends the filter balance as implemented starting in `baseline2.0` should not be changed, it is possible that rebalancing the exposure time to compensate for performance and throughput in some filters as compared to others or shortening exposures in filters where the throughput exceeds expectations enabling the collection of more images in that filter (or overall) would lead to enhanced LSST science. The SCOC cannot finalize this recommendation at this time due to missing information about the characteristics of the system-as-built.”

Simulations of the survey strategy up to and including `baseline_v3.2` use throughput curves assuming mirror coating as Al-Ag-Al respectively for M1-M2-M3. Conversely, the plan to coat the mirrors was updated in 2023 to Ag-Ag-Ag (or 3xAg), which leads to a $\sim 15 - 20\%$ increase in survey efficiency compared to Al-Ag-Al by increasing throughput in all bands redder than u , and bringing throughput closer to the design goals as stated in SRD.¹ Table 1 shows the magnitude limit changes associated with coating. 3xAg coating increases sensitivity in all LSST bands from g to y , but it decreases the throughput in u . As of `baseline_v3.3`, all OpSim simulations include the updated expected throughput.

	Al-Ag-Al E2V	Al-Ag-Al ITL	Ag-Ag-Ag E2V	Ag-Ag-Ag ITL
u	0.0	-0.06	-0.21	-0.27
g	0.0	-0.04	0.06	0.02
r	0.0	-0.05	0.10	0.05
i	0.0	-0.02	0.13	0.12
z	0.0	0.01	0.15	0.15
y	0.0	0.03	0.07	0.10

¹<https://community.lsst.org/t/rubin-sim-v1-3-released/7937> and https://github.com/lsst-pst/syseng_throughputs/blob/main/notebooks/SilverVsAluminum.ipynb.

Table 1: Magnitude limit change, compared to expected magnitude limit with E2V chips coated with Al-Ag-Al. Note: values are reported for chips acquired by two different filter vendors as E2V and ITL for both coating schemes. Positive values indicate deeper limiting magnitudes.

The SCOC reviewed the largely positive impact of the new throughput on science cases: nearly all MAFs responded positively to the increase in survey depth with some metrics showing improvements as large as 10% (e.g. Parallax uncertainty, see Figure 4, see also Figure 1 and note the significant improvements between baseline_v3.2 and baseline_v3.3 when the new filter transmission curves were introduced).




FIGURE 4: Gains in the metric tracking LSST's median parallax uncertainty at magnitude $r = 24$, a standard metric that measures a LSST SRD system requirement for different baseline OpSimS, from baseline_v1.x through baseline_v3.3, the first simulation with updated system throughput reflectivity from the 3xAg mirror coating. The improvements in parallax uncertainty between v3.2 and v3.3 OpSim come from the change in throughput and added depth in all bands bluer than u . Similar improvements are seen in proper motion uncertainty.

However, the SCOC understands that the throughput loss in u -band (~ 30% loss in coated)

depth) would negatively impact science cases including Photo- z , studies of the Milky Way halo, Lyman Break Galaxies (LBGs, identified as u -band dropouts at redshift $z \sim 3$), and more. While the overwhelming majority of the MAF metrics available to the SCOC responded positively to the updated throughput, we are aware, as always, that these may not provide an exhaustive picture of the science outcomes. Therefore, guided by experts in the community, we explored ways to reduce the u band magnitude decrease while preserving the benefit of increased throughput in redder bands. We tracked the performance of Photo- z , as characterized in Graham et al. 2017, by assessing the variance and bias in Photo- z at redshifts $z \lesssim 3$. Photo- z is sensitive to u band depth at redshift $z \geq 2$ due to decrease power to identify Lyman break galaxies photometrically. We expect that recovering Photo- z performance is a good indicator of recovering performance for other science cases sensitive to u band depth for which we do not have detailed metrics. Photo- z performance, along with a large set of MAFs, were run against a set of `OpSims` that progressively changed the exposure time and the number of exposures in u -band².

The SCOC recommends:

- **an increase of the exposure time in u -band to 38 seconds per visit**
- **an increase of the number of u -band visits of 10% compared to `baseline_v3.0`**
- **an identical decrease of 0.8 second exposure time in all other bands to compensate for the added time in u -band.**

This roughly restores the u band depth of LSST in `baseline_v3.0` with minimal impact on LSST science metrics. As a science case that is representative of those sensitive to u -band depth, these changes recover performance on Photo- z at redshift $z \sim 2$, where the impact of the u -band throughput loss was most significant while maintaining the performance improvement on Photo- z at low redshift afforded by the increased depth of LSST in all other bands. Furthermore, these changes minimally impact other science cases tracked by MAFs.

Since more science cases generally respond better to increasing the number of images, over increasing the exposure time to achieve the same depth, the added u -band time should be obtained by decreasing (minimally) the exposure time in other bands, rather than decreasing the number of images.

²<https://community.lsst.org/t/release-of-v3-4-simulations/8548>

We note that this recommendation is subject to ongoing feasibility studies by the Rubin Data Management team.

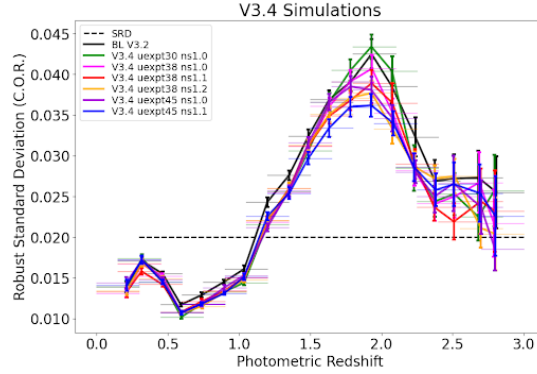


FIGURE 5: Effect of changes of u -band 10-year depth on Photo- z Robust Standard deviation as a function of redshift z as measured in Graham et al. (2017). The dashed line represents the SRD requirements on Photo- z . The black solid curve is the `baseline_v3.2`, the latest baseline before the filter transmission curves were updated in the Rubin simulation system, the colored curves represent the Photo- z Robust Standard Deviation varying u -band exposure time between $30 \leq u_{\text{expt}} \leq 45$ seconds and number of exposures in u band between $1.0 \times ns \leq N_u \leq 1.1 \times ns$ where ns is the number of u -band exposures in `baseline_v3.2`. Note that, with the caveat that sampling uncertainties are large at $z > 1.5$, we note an improvement in Photo- z at low z associated with increased depth in bands redder than u , but a degradation at $z > 1.5$. A similar impact is seen in Photo- z bias.




FIGURE 6: A standard set of science and system MAFs metrics as a function of changing exposure time ($27 \leq u_{expt} \leq 45$ seconds) and fraction of exposures in u band ($0.9 \times ns \leq N_u \leq 1.2 \times ns$). The metrics are normalized with respect to a simulation with $u_{expt} = 30$ seconds and $N_u = 1.0 \times ns$. Three additional columns on the left show: $v_{3.2}$ (baseline_v3.2, pre-filter-throughput update, notably generally worse) and $v_{3.3}$ (baseline_v3.3, which follows the same observing strategy as to baseline_v3.2 but includes throughput updates) and $u_{38s} 1^*$ where the exposure time of all other bands is adjusted to compensate for extra time spent in u (whereas in all other simulations shown in this plot the exposure time is kept at 30 seconds). The SCOC indeeds recommends an adjustment of the exposure in all bands (29.2 seconds instead of 30 seconds) and this is implemented in all simulations starting with $v_{3.5}$.

3.3 Rolling

In a rolling strategy, instead of distributing visits uniformly on the WFD footprint, the sky is split into regions that alternate high- and low-intensity monitoring. In PSTN-055, the SCOC recommended the implementation of a rolling strategy for the LSST WFD at a strength of 0.9³ with the sky split into two rolling regions constituted by four longitudinal stripes. The primary drivers for this recommendation are time-domain science including the exploration of the transient and variable sky and SN cosmology. Rolling as described decreases the median time-gaps compared to a no-rolling implementation of LSST: distributing the ~800 visits per pointing evenly into 10 seasons results in an median revisit time per pointing of about 4.5 nights, while rolling can increase the cadence on the areas of sky closer to 2.5 nights.⁴

However, concerns were raised by the DESC, and seconded by other groups such as the Galaxies SC, regarding the lack of uniformity in the distribution of depths across the survey at intermediate years between 1 and 10 compared to a no-rolling strategy (see discussion in Lochner et al., 2022, which noted this as a potential future concern before the adoption of a rolling cadence as the baseline). These concerns highlighted the negative impact that rolling induces on the cosmological analysis conducted with static-sky probes due to a decreased uniformity of the data releases, which has been shown to cause several significant issues for cosmological large-scale structure analyses (e.g., Abbott et al., 2022; Baleato Lizancos & White, 2023). This uniformity challenge could be addressed after data collection by selecting and limiting the number of images going into a data release per field (including these data in future data releases) to achieve higher uniformity or by a “renoising” step. These data management solutions are at the moment unscoped and do not fall under the current requirements of the Rubin DM deliverables.

[PSTN-055] “The current SCOC recommendation is to implement a rolling cadence with a half-sky rolling scheme and a 0.9 rolling weight. However, rolling impacts the uniformity of static data releases which, as experts in the community have highlighted, is necessary for static sky science in general and cosmology in particular. This issue may be resolved or mitigated at the software level in the creation

³This number represents the fraction of the visits that the scheduler attempts to place in the high-activity rolling region. However, the resulting visit distribution is more uniform (75-80% in high-activity regions, 25-20% in low ones) due to competing requirements (e.g., filter balance, minimum number of observations per pointing per year in each filter to produce templates, weather, etc...)

⁴The reader is reminded that each pointing receives two or three visits per night. The time gaps reported here are for inter-night observations.

of coadds and catalogs, rather than at the scheduler level. The community should specify the desired and necessary requirements for uniformity to enable the exploration of data processing solutions to this problem. Depending on the feasibility of a solution to ensure sufficient uniformity, the SCOC recommendation on rolling may be re-evaluated.”

With the goal of quantifying the necessary uniformity to enable cosmological results at certain key data releases⁵, DR5 and DR8, as well as identifying solutions that enable rolling (at the strength recommended in PSTN-055) while increasing the uniformity of key data releases, a Uniformity Task Force developed alternative rolling implementations. *Uniform rolling*, implements interruptions of rolling before specific data releases to increase the uniformity of those releases and recover an acceptable level of uniformity at key years – see Figs. 8 and 9. It was found that uniform rolling permits the full survey area to be used for cosmological analysis at years 4 and 7, whereas in previous rolling versions, approximately 35% of the cosmological constraining power⁶ was lost at those years due to the need for area cuts.

We note that rolling can be implemented with at most four rolling cycles (that is, starting rolling in Y2 and ending rolling in Y10) where a cycle is defined as a pair of two years where the high- and low-intensity regions are swapped. Uniform rolling requires limiting rolling to three cycles. Primarily this affects science sensitive to timescales of $\sim 24\text{--}48$ hours⁷. We note that these time scales had been identified as sensitive and requiring additional improvements in PSTN-055 within the recommendation on rolling (PSTN-055 Section 2.4.1):

[PSTN-055] “[...] the SCOC recommends the LSST cadence be designed to ensure coverage of time scales in the hours-to-one-day range by carefully tuning survey parameters in combination. Performing three visits per night by default is not recommended, but a combination of preferentially pushing a third visit to the follow-

⁵These intermediate releases were selected because they enable equally-spaced time intervals between new datasets for comprehensive static science analysis: years 1, 4, 7, and 10 corresponding to DR2, DR5, DR8, DR11.

⁶Here we quantify cosmological constraining power through emulated forecasts of combined constraints from cosmological weak lensing and large-scale structure measurements (Lochner et al., 2022). The constraints assume a w_0w_a CDM cosmological model, with w_0 and w_a entering as two parameters in the dark energy equation of state. The constraining power is quantified through the area of the uncertainty contours in the (w_0, w_a) part of parameter space, marginalizing over other cosmological parameters and systematic uncertainties – then taking the inverse of that area (so that higher values mean lower uncertainty, i.e., tighter cosmological constraints). However, this can be considered more generally as a proxy for how well we are measuring cosmological structure growth, translating into tighter constraints on the amplitude of matter fluctuations if a Λ CDM cosmological model is assumed.

⁷Shorter time scales are primarily covered by observations in triplets as discussed in PSTN-055.

ing night [...] and requesting a third visit within a night once every several nights (~ 1 week) would achieve this goal.”

While not designed for this purpose, the implementation of the 0.9 strength, two-sky-areas rolling with four cycles (baseline_v3.2 and later) improved coverage at 24-48 hours (see figure Figure 7).




FIGURE 7: The mean number of observations at 24, 48, 72, and 96 hours in all bands, *g* band only, *r* band only, and *z* band only as a function of the number of rolling cycles: each cycle of rolling cadence improves the number of samples by about 7%. On the left, the ratio of samples normalized to the number of samples when not rolling is shown for two OpSims with 3 rolling cycles (roll_3_v3.4_10yrs and roll_uniform_early_half_mjdp0_v3.4_10yrs) and two OpSims with 4 rolling cycles (baseline_v3.3 and baseline_v3.4). On the right, the absolute number of samples is shown in *g*, *r*, and *i*. Note the overall small numbers of samples in these time scales when not rolling: < 5 in *g* and $10 < ns < 15$ in *r*, *i*, *z* at 24 hours. The rolling strategy improves sampling between 24 and 96 hours, but recall that PSTN-055 concluded that the sampling enabled by the rolling as implemented in baseline_v3.0 (in three cycles) was still insufficient. A fourth cycle of rolling improves this sampling by ~ 7%.

The SCOC recognizes the positive impact this rolling implementation has on static cosmological and extragalactic probes and considers this a promising solution for the uniformity concerns raised in PSTN-055, with limited impact on time-domain probes. However, this implementation of rolling is a significant and new departure from earlier implementations, and the number of cycles of rolling had not been previously explicitly discussed as a parameter in the survey strategy. Furthermore, the current Uniform Rolling implementation requires rolling to start early in Y2 (in the current implementations rolling starts on survey day ≤ 400) but, as discussed in PSTN-055, rolling shall not start until sufficient sky coverage has been achieved to enable proper photometric calibration.

For these reasons, the SCOC is not committing at this time to recommend any specific im-

plementation of rolling, beyond confirming the strength 0.9 and two-region strategy. Since in all current implementations rolling does not begin until Y2, the SCOC intends to continue investigating rolling implementations and their impact throughout Y1, with the support of the community, and release a recommendation of how to implement rolling in its first annual recommendation ahead of Y2 of Operations. In particular, we intend to (1) investigate sensitivity to the outcomes of Y1, (2) ensure the community has time to evaluate the potential impacts of these changes that are not currently highlighted by our metrics, and (3) refine the uniform rolling implementation details.

The SCOC recommends that the time domain community, particularly those interested in phenomena that have evolutionary time scales of hours-to-days, urgently quantifies the impact of the proposed uniform rolling compared to rolling in four cycles. For this purpose, while the baseline is implemented with 3-cycles uniform rolling, the Survey Strategy team has prepared v3.5 OpSims with different rolling implementations.

Further, the SCOC iterates the recommendation that Data Management scopes a plan for producing uniform data releases in DR5 and DR8 in addition to the standard data releases and that the cost of the development and storage of these additional data releases is scoped and shared with the scientific community. Even if produced by the Rubin DM, uniform data releases will require the input (of the DESC and extragalactic science community at large) to develop the algorithm that will achieve sufficient uniformity. Understanding the cost of producing two additional *uniform* data releases is necessary to compare this cost to the scientific cost of three vs four cycles of rolling, to be measured by the community (see previous paragraph). In addition, if rolling cannot start early enough to interrupt rolling ahead of DR5 and DR8, this remains the only alternative solution currently identified to achieve sufficient uniformity. Sharing information on the cost of additional data releases will place the community in a position, if needed, to advocate and secure funding for this purpose.

The SCOC is thankful to the Uniformity Task Force, chaired by Rachel Mandelbaum, which provided invaluable contributions and analysis that led us to this recommendation.

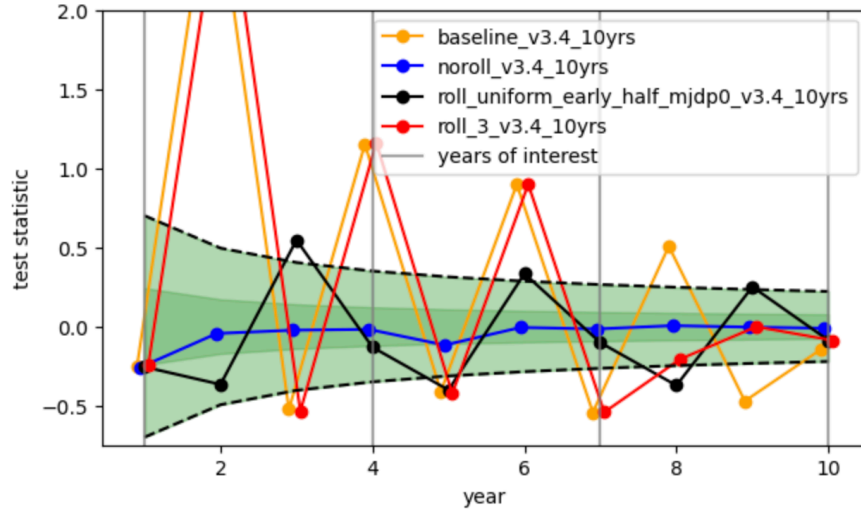


FIGURE 8: Quantitative assessment of the non-uniform exposure time variation vs. year under different observing strategies. The test statistic plotted on the vertical axis effectively measures the fractional difference between the variations in depth between the northern and southern Galactic regions, with a value of 0 indicating that the two are the same, as expected for a perfectly uniform survey. The light green shaded envelope between the dashed black lines indicates the region for which we consider the stripe features to be negligible (meaning manageable within the limits of existing analysis algorithms). The narrower dark green shaded envelope shows the expected statistical fluctuations for a survey without rolling, as estimated using the `noroll_v3.4` strategy simulation. As shown, at the highlighted years (1, 4, 7, 10) the uniform rolling strategy (`roll_uniform_early_half_mjdp0_v3.4_10yrs`) is very close to uniform within the level of statistical fluctuations at Y4 and Y7 (DR5 and DR8), while the `baseline_v3.4` strategy is highly non-uniform, especially in those years.

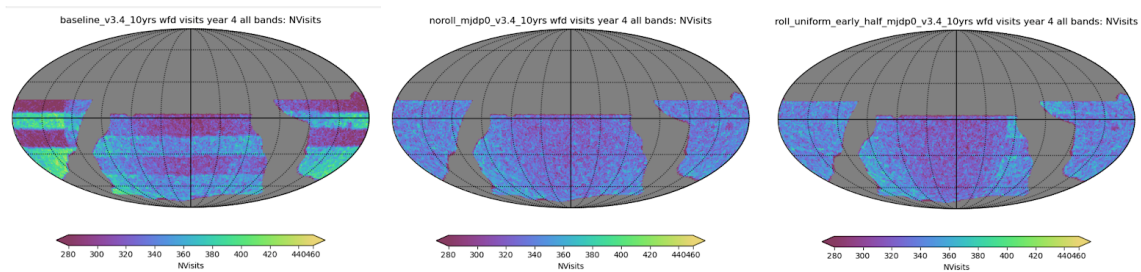


FIGURE 9: Comparison of the depth of LSST at the end of Y4 (DR5) under different rolling strategies. The left plot shows the LSST number of visits map for a standard implementation of rolling at strength 0.9 in two sky regions designed as four longitudinal stripes. The center plot represents an implementation of LSST without rolling, for comparison and provides an upper limit to expected uniformity. The right plot shows the implementation of Uniform Rolling described in this section and implemented in `baseline_v4.0`.

3.4 Galaxy

In PSTN-055 the SCOC identified areas of work needed to finalize the WFD survey strategy on the Galactic sky and special regions of interest to Galaxy science, including the Clouds and South Celestial Pole, which can be observed within the WFD but with different observing choices than the low-dust footprint, of primary interest for extragalactic science. This section includes updated recommendations on the Galactic footprint and its observing cadence, including whether rolling should be implemented (subsubsection 3.4.1), filter balance (subsubsection 3.4.2):

[PSTN-055] “The SCOC is not ready to finalize a recommendation for the filter balance in the Galactic Plane, or for a final Galactic Plane/Bulge footprint, or the rolling scheme to be implemented on the Galactic Plane. The SCOC will work with the SMWLV and TVS SCs to ascertain the best solutions for Galactic science on filter balance and footprint. These decisions should, however, not impact decisions relating to the WFD and the time spent collectively on Galactic regions should not change. Galactic Plane pencil-beam surveys need to be defined more clearly to assess if they would ultimately result in “nano-surveys”, which will require a fraction of time too small to be optimized at this stage, or to evaluate the possibility of incorporating them in a final Galactic Footprint recommendation”.

and special regions (subsubsection 3.4.3).

3.4.1 Footprint and Time Distribution of Visits

Extensive work has already led to the present division of the dense regions of the Galaxy into a high-visit region that encompasses both a large area around the Bulge and a long, thick strip of the Plane, surrounded by a larger area in the Plane with fewer visits. The subsequent efforts of the SCOC and scientific community have been focused on refining these choices.

One feature of the `baseline_v3.0` (PSTN-055) survey in the Plane is that it left a high-visit “blob” centered around a Galactic longitude of $l = +45$ surrounded by a lower visit area. This resulted from a previous candidate survey design that included high-visit pencil beams⁸ at varying

⁸In a subset of previous candidate survey designs, “pencil beams” were a series of 20 high-visit single pointings distributed in galactic longitude with the goal of ensuring the survey sampled a range of stellar environments.

Galactic longitudes along the Plane and considered stellar density, but was not due to any other specific science goal in this region. Visits centered around this high-declination blob would necessarily have to occur at high airmass, and would additionally be separated from other high-visit areas, reducing survey efficiency.

The SCOC recommends redistributing the excess visits in the “blob” centered around a Galactic longitude of $l = +45$ to cover a low-visit “barrier” at $l = +335$ in the Plane and at the border of the Plane and Bulge. This change would give continuous longitude coverage along the Plane from a longitude of $l = +30$ down through $l = +280$ and boost metrics for time-domain science in the Bulge/Plane.

The SCOC recommends rolling on the low-dust WFD (see subsection 3.3), where strips in declination alternate high- and low-intensity monitoring. However, this rolling implementation need not extend to the dense regions of the Galaxy, if it does not provide overall scientific benefits to Galactic science.

In `baseline_v3.0`, no rolling is implemented in the Bulge and Plane footprint. The SCOC explored simulations that implemented rolling in both regions or in only the Bulge. Rolling in both the Bulge and Plane is extremely unfavorable for many Galactic transient metrics, such as microlensing discovery and characterization for a broad range of event timescales, as well as early detection of X-ray binary outbursts.⁹ The outcomes are more complex for rolling in the Bulge alone; while still negative for Galactic transient discovery, rolling in the Bulge has a mixed effect on microlensing metrics. These Bulge simulations particularly aimed to explore whether rolling cadence implementations could boost the early detection and characterization of shorter (\sim few days) timescale microlensing events and anomalies, even if only for a limited survey region. In practice, the improvement was found to be comparatively small, and came at the detriment of the regular, long-baseline monitoring necessary to characterize long-timescale events such as those caused by compact object lenses.

The SCOC concluded that rolling on the Galactic footprint would have a net negative effect on the survey as a whole, and recommends no rolling in the Plane or Bulge.

Finally, the SCOC recommends redistribution of a small number of Bulge visits to a central Bulge field overlapping the planned Roman Bulge survey area, with a goal of more contin-

⁹The X-ray binary outburst metric is representative of Galactic transients with a typical duration longer than a few days that follow the stellar distribution in the Galaxy, so has much broader relevance than solely for X-ray binaries.

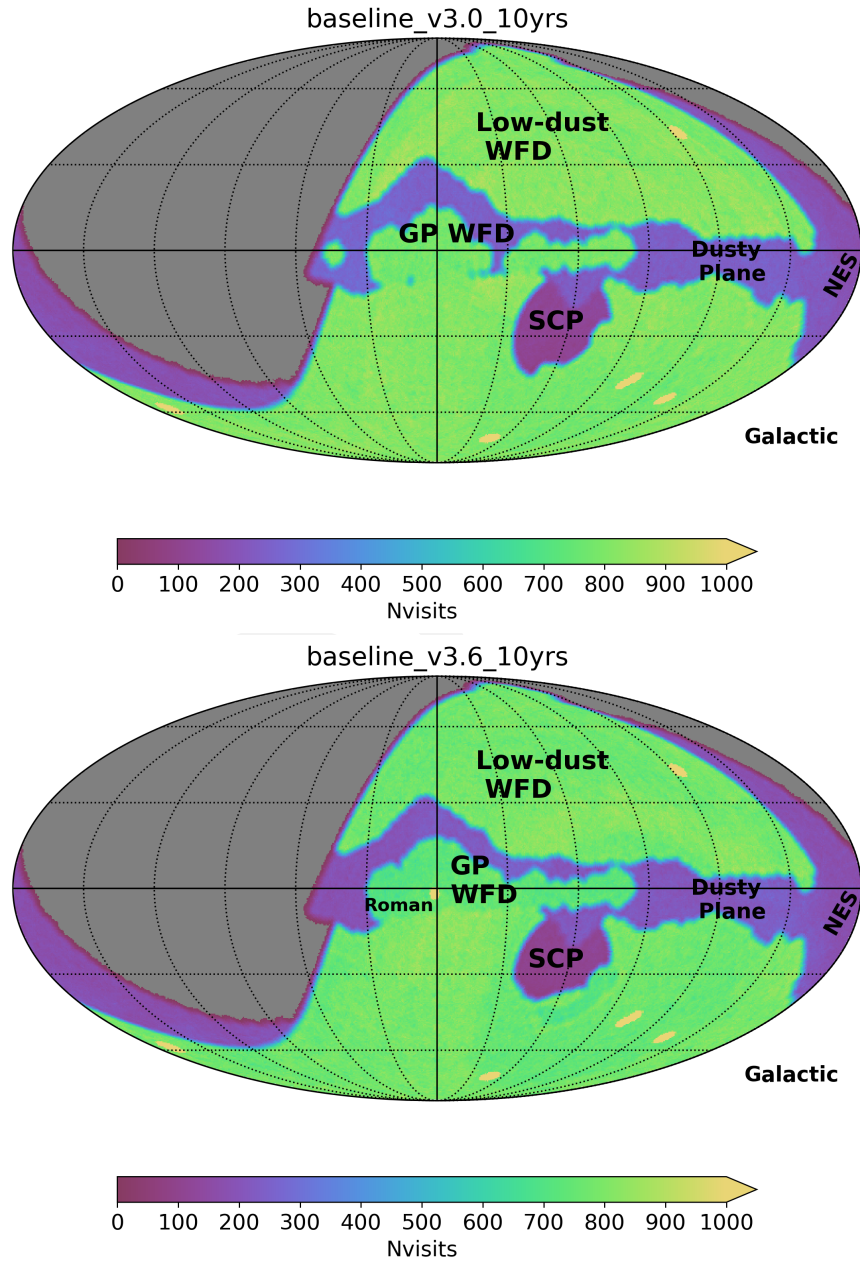


FIGURE 10: Comparison of the MW footprint (Galactic coordinates) as recommended in PSTN-055 and implemented in `baseline_v3.0` and through v3.5 simulations and the refined footprint recommended in this document and implemented in `baseline_v3.6` and later.

ual monitoring to improve microlensing detection and characterization in this region that will be intensively surveyed by Roman during predetermined seasons. This recommendation was first implemented in `OpSim roman_v3.3` and resulted in large improvements for some microlensing metrics, improvements for all microlensing metrics, and no significant negative impact on any other metrics.

The SCOC recommends that a visit plan consistent with this `roman_v3.3` with the number of redistributed visits to be capped at that used in this simulation: in the current Roman field `OpSims` the number of visit reallocated to the Roman field is ~ 1600 . However, the timing of the implementation of this augmented observing campaign needs to remain flexible at this time to respond to the as-of-yet to be finalized launch date of Roman and the scheduling of its surveys.

3.4.2 Galactic Filter Balance

The filter balance in the Bulge in the current `baseline_v3.4` already differs from that used for WFD: the primary difference is fewer visits in y , which are redistributed to bluer filters to better optimize Galactic science since u and g are vital for stellar characterization even in the presence of foreground dust. In WFD y receives a large number of visits, comparable to z and only slightly less than r or i , while g and especially u receive fewer visits. Hence y , with its relatively low sensitivity, is the optimal choice for redistribution to bluer bands.

Noting the relatively low sensitivity of the y band, and its resulting negligible reddening advantage over z even in dusty regions, the SCOC considered several simulations that redistributed additional visits in the dense regions of the Galaxy from y to a combination of z , g , and u , while still recognizing the fundamental discovery potential of a multi-filter survey over a broad contiguous area. The main finding from these new simulations was that most existing metrics showed mixed or marginal changes, even where the relative number of visits in u and g substantially increased. The metrics considered included Galactic transients, young stars, detection of several classes of periodic variables, light curve gaps, as well as solar system metrics (since the ecliptic passes through this region).

The SCOC finds that the adoption of a revised filter balance in the Bulge and Plane with less y and more z , g , and u compared to the present baseline is potentially beneficial on the net, but that existing metrics are not adequately sensitive to the explored filter balance changes for some expected science cases. The SCOC concludes that a survey

using the currently implemented filter balance in the Bulge and Plane in baseline_v3.4 will produce excellent science and the LSST can start with this implementation. However, the SCOC also welcomes input from the community whose science is affected by the details of filter balance in the dense regions of the Galaxy to help define improved metrics that could lead to further optimization in future years.

3.4.3 The Milky Clouds and South Celestial Pole

The scientific goals of the survey in the region of the Large and Small Milky Clouds (LMC and SMC, together MCs)¹⁰ and South Celestial Pole (SCP) differ somewhat from WFD. In particular, the major areas of focus of the survey of the main bodies of the LMC and SMC are microlensing and other variable/transient science. In the peripheries of the MCs, including the South Celestial Pole region, the central goal is to detect dwarf satellites and other low surface brightness stellar substructures. These goals are supported in the current baseline, as the MCs are covered with visit numbers like that in WFD, while the SCP region, only observable at relatively high airmass, has a low number of total visits, but sufficient to detect many potential dwarf satellites and substructures. However, the current baseline also adopts the WFD filter baseline in the MC and SCP regions, which may not be ideal for the stated goals.

A number of simulations were considered that used an alternate filter balance for both the Clouds and the South Celestial Pole, moving visits out of z/y and toward u/g in both regions. These simulations show large improvements in metrics relevant to the detection of low surface-brightness dwarfs as well as some improvements in microlensing and variable star/transient metrics.

The SCOC recommends a bluer filter mix in these regions, bounded by the requirement that the increased number of dark-time visits in a relatively narrow range of right ascension does not affect other parts of the survey.

The SCOC is thankful to the Galaxy Survey Strategy Task Force, chaired by Jay Strader and Rachel Street, which provided invaluable contributions and analysis that led us to this recommendation.

¹⁰There is an effort underway to have the community avoid using the name Magellanic Clouds, as reasoned in <https://physics.aps.org/articles/v16/152>. Milky Clouds has been proposed as an alternative and has indeed been used in a number of settings. We adopt this name here to reflect the broad and inclusive reach of Rubin LSST.




FIGURE 11: The different filter balance in the LMC and SMC region compared to the rest of the WFD can be seen when comparing the number of exposures in g , and i , and z at the end of the 10 year survey (the rectangular region near the SCP). These figures also demonstrate the bluer filter balance in the Milky Way region.

3.5 Targets of Opportunity (ToO)

In PSTN-055 the SCOC recommended the implementatino of a ToO program that should:

[PSTN-055] “be contained to $\leq 3\%$ of the LSST time. The SCOC recommends that Rubin organizes a workshop in 2023 to bring together members of the scientific community, members of Rubin Observatory (including observing and scheduler specialists, and Data Management specialists) and members of the SCOC to define the details of the implementation of the Rubin ToO program. This workshop should produce a document detailing recommendations for implementation, including suggestions for the questions outlined above, that the experts agree would accomplish the scientific goals of the program.”

A meeting was organized in March 2024 (Rubin ToO 2024¹¹) with the explicit purpose of making a community recommendation for a Rubin ToO program within bounds previously established by the SCOC. After evaluating this community consensus report and considering simulations of its implementation, the SCOC finds that the impact on WFD science is generally small and that the proposed ToO programs have the potential to lead to important scientific results.

The SCOC recommends the implementation of a LSST ToO program as detailed in the community report Rubin ToO 2024: Envisioning the Vera C. Rubin Observatory LSST Target of Opportunity program ¹² (hereafter *RubinToO2024*) by the scientific community at large.

¹¹https://lssttooworkshop.github.io/images/Rubin_2024_ToO_workshop_final_report.pdf

¹²https://docs.google.com/document/d/1WE4NGl3dFOVGo7lZpyG1fe_JiX9m-kL15JYQkhu9iso/edit?usp=sharing

The impact of including a ToO program as recommended in RubinToO2024 is shown in Figure 12. In the current implementation, the program takes between 3% and 4% of the survey time. While this is slightly in excess of the recommendation in PSTN-055, we are still improving the efficiency of the program's implementation, and the current implementation likely represent an upper limit as no triggered sequence is terminated due to reclassification of the event and/or as the transient is identified.

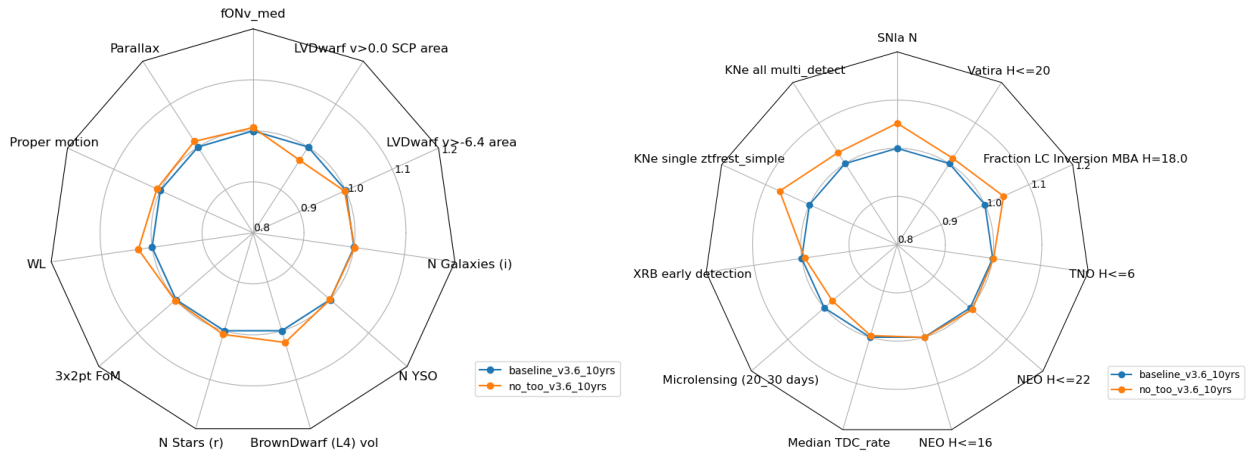


FIGURE 12: The impact of the inclusion of the ToO program on static (left) and transient and variable (right) LSST science. Note that the marginal negative impact on the number of well-characterized SNIa (SNIa N MAF on the right) and Kilonova discovery (KN- MAFs in the same plot) in the WFD is compensated, respectively, by the potential for discovery of KN counterparts of MMA triggers, and by the promise of KN counterparts of gravitational waves as cosmological probes (e.g., Coughlin et al., 2020; Gianfagna et al., 2024).

This report identified several different classes of ToOs for which Rubin's observations are well-justified. The vast majority of ToOs will be to follow up gravitational wave (GW) events, while a much smaller number of neutrino and solar system ToOs are expected.

The current LIGO-Virgo-KAGRA (LVK) GW observing run (Observing Run 4 or O4) will end before the start of LSST. Hence, GW ToOs will not commence until the start of the Observing Run 5 (O5) of the LVK detectors. We note that the start time of O5 has no expected impacts on the LSST WFD or the ToO program. Improved system performance, primarily afforded by the consistent working of three detectors (with similar sensitivity), will maximize the scientific productivity of the Rubin ToO program while reducing the impact on other programs. Two working LIGO detectors at their design sensitivity, combined with a third detector working at 30-50% that of LIGO, will reduce the skymaps to tractable sizes for rapid Rubin coverage. We encourage the LVK science collaboration and the International Gravitational Wave Network (IGWN), to prioritize a high-performing system with three working detectors over an early

start of the O5 run.

As the GW component of the ToO program takes the largest amount of time and has the most impact on WFD, to enable optimal use of Rubin resources, **the SCOC recommends that a meeting to follow Rubin ToO 2024 is organized closer to the start of O5 to refine the GW follow-up survey strategy with improved knowledge of the expected performance of the GW detector networks and systems in O5 and of the performance of the full Rubin system.**

There is no comparable time restriction for the solar system ToO program (to follow up potentially hazardous asteroids) or neutrino ToO program (to follow-up high energy neutrinos or those from a Galactic supernova). Hence, **the SCOC recommends that the solar system and neutrino ToOs should start as soon as possible: as soon as suitable templates are available.**

For all ToOs, to enable ToO response from the Rubin system, a high level of automation is required. For each potential ToO, a response shall be pre-determined algorithmically, including which targets Rubin responds to and the sequence of observations, based on the transient's characteristics. Informal systems can easily lead to mistakes. For this reason, **the SCOC recommends that Rubin only consider potential ToOs that emanate from vetted discovery and distribution systems that produce and dispatch fully machine-readable alerts.**¹³ The SCOC considers the current list of vetted systems to be: LIGO-Virgo-KAGRA (gravitational waves); IceCube (neutrinos); SNEWS (neutrinos); JPL Scout or Sentry for potentially hazardous asteroids. The SCOC will evaluate future systems for inclusion in this list (e.g., a new neutrino observatory) on formal request.

Human input may still be required to evaluate in real-time the value of a ToO trigger and the specific response. One (or more) Rubin members should review triggers and be allowed to, if desired, overwrite the algorithmic decision to pursue/not pursue a ToO or interrupt the ToO observing sequence. Further, to ensure that appropriate expertise is available, an Advisory Committee of community members can interact and advise the observer in real-time, with communication initiated either by the committee or by the observer. This Advisory Committee should be composed of community members and have a nomination-selection process (including self-nomination) to be outlined in detail before the start of survey operations, en-

¹³At the time of writing, the SCOC understands that full automation is not currently in place for all IceCube neutrino triggers.

surveying broad coverage of scientific competence and diversity along all relevant axes. The committee, observers, and Rubin leadership will review the ToO outcomes post-factum to advise on program changes. **The SCOC recommends real-time human review of potential ToO triggers and the establishment of a Rubin ToO Advisory Committee as described above.**

3.6 Snaps

While the LSST has originally been designed to collect two 15-second snaps for each visit, primarily to remove cosmic rays to be combined into a single visit that is subsequently processed to derive the data products, there is an opportunity to move to collect a single 30-second exposure¹⁴, as it appears that cosmic rays can be reliably rejected from a single image. The feasibility of this plan remains to be ascertained in commissioning (including from on-sky images). However, the SCOC has conducted a science-driven analysis of this proposal.

Going from 2x15s to a single 30s brings a gain in efficiency equivalent to 7-9% of the survey time (associated with reduced read-time).

Saturation limits will be slightly higher but this will only impact a small number of objects compared to the large volume of sources in the LSST universe. Other surveys are better equipped to work with those targets that are bright for LSST. Some science cases (e.g. Cataclysmic Variables and Flares, very fast moving solar system objects) could benefit from the separate exposures, but data products associated with the snaps would not be released individually to the public in Data Releases, only the combined image and flux difference where detected are official data products, so all those science cases would depend on pipelines contributed by the community and user-generated data products. Additionally, for these cases too, other surveys are better equipped to work within those time scales.

Thus, the SCOC does not see scientific opportunities associated with retaining the two 15s snaps that can compete with the 7-9% gain in survey efficiency.

The SCOC recommends that, if the technical feasibility is confirmed in commissioning, the survey is conducted with single exposures. With our recommendation of modifying the exposure time for u band to 38 seconds, and compensating for this extra u band

¹⁴Note that joint with the recommendation on u band exposure length and filter balance (subsection 3.2) the total exposure time in all bands but u drops to 29.8 seconds, and the exposure time in u is 38 seconds. However, for convenience, we will continue to refer to a "2x15 seconds" and "1x30 seconds" implementation.

survey time by a short decrease in exposure across all other bands, the single visits would be 1x29.2 seconds.

The time gained by avoiding snaps will not be allocated to any specific program in Y1 as the performance of the system is still uncertain. In the future, the SCOC will consider how the additional time may be allocated, including to special programs (e.g., microsurveys), DDFs, WFD, etc., either to compensate for unexpected performance loss or to increase science throughput.

3.7 Deep Drilling Fields (DDF)

A general plan for the LSST Deep Drilling Fields (DDF) has been developed over the course of the past 15 years, starting with LSST Science Collaboration (2009) through many further developments and recommendations (e.g., Brandt et al., 2018; Scolnic et al., 2018). The DDF program will include 5 DDF pointings. The SCOC recommended in PSTN-055 that 6-7% of overall survey time is dedicated to the DDF program, and that each DDF receives 20k visits except for the COSMOS field which should receive 40k (with accelerated coverage so that COSMOS reaches 20k visits by the end of Y3), and the The Euclid Deep Field South (EDFS) has a wider area equivalent to two separate pointings (sharing 20k visits across the two pointings). However, some questions remain:

[PSTN-055] "The SCOC will continue working in 2023 with the community to identify the specific intra-night cadence that maximizes the science throughput of the DDF survey, while not impacting the science performed by other surveys."

[PROVISIONAL RECOMMENDATION] The implementation of DDF intranight visits is still under development, retaining sequences of visits in multiple filters within a night, providing opportunities for fainter per-night measurements as well as sub-minute sampling for limited times. Tradeoffs between nightly depth, cadence, season length, and filter balance are still being explored, but the SCOC is able to make some further recommendations based on input from the SCOC DDF Task Force.

The SCOC recommends that DDF observations should be sequences of multiple WFD-like visits (as opposed to increased exposure times) to allow rapid alert generation. This also benefits ~minute timescale science and cross-calibration of DDF and WFD observations.

The SCOC recommends that the baseline translational dithering scale of DDF observations should be reduced from 0.7 degrees to 0.2 degrees (with exploration of even smaller translational dithers compatible with instrumental signature removal and calibration needs).

Smaller translational dithers allow DDFs to reach increased co-added depth for static science and increased temporal coverage for time-domain sources. While larger dither is favored for low-surface brightness science, no clusters or low-surface brightness structures of interest are included in the DDF fields.

The SCOC recommends that the baseline survey strategy should accommodate varying the nightly depth, filters, or cadence of different DDFs throughout the course of LSST, while maintaining the Phase 2 recommendations for the 10-year depth of each field (including the enhanced COSMOS observations to reach 10-year depth in the first 3 years).

Adding this flexibility to DDF observations allows for periods of higher cadence necessary for some transient science (*e.g.*, AGN or supernovae) and enables more opportunities for concentrated, contemporaneous observations with other surveys (*e.g.*, Euclid, Roman) while maintaining the overall co-added depth for static science.

The SCOC urges the Data Management and Alert Production teams to assess the feasibility of, and resources needed for, enabling nightly co-adds of sequential DDF visits and recommends that a path is developed to enable the creation of these co-adds, subtraction with deep templates, and faint alert generation (with higher latency as needed, *e.g.*, after sunrise).

Nightly co-adds are required to take advantage of the increased DDF depth in the time domain. Alerts from nightly co-adds are essential for faint time-domain sources (*e.g.*, high-redshift AGN or supernovae). Longer timescale co-adds (*e.g.*, weekly, monthly, yearly) and alerts should also be considered.

3.8 Early Science

The SCOC emphasizes that the priority in Y1 of operations should be obtaining a dataset that supports and facilitates science throughout the survey. This includes a dataset sufficient for

calibration across the $\sim 20,000$ square degrees of the WFD, including images at different air-masses, illuminations, field crowdedness etc.

The SCOC supports Rubin's commitment to acquiring incremental templates throughout Y1 to begin dispatching alerts (via the Alert Brokers) and encourages the Observatory to release alerts as early as possible. The SCOC reviewed the Alert Production team proposal to prioritize timeliness over the quality of templates and build templates from fewer images (≥ 3) in Y1 than in subsequent years. Releasing some alerts in Y1 is an important goal to enable the time domain science community to prepare for the full-volume full-fidelity alert streams to come in subsequent years, as well as increasing the discovery potential of LSST in early operations. However, this goal should not overwrite the priority of obtaining a fully calibrated system by the end of Y1.

The SCOC recommends that the filter balance is slightly modified in Y1 to acquire a sufficient number of u band images for calibration.

The SCOC does not recommend beginning rolling before the end of Y1. In the months following the release of this recommendation, the SCOC will continue to work on the implementation of rolling (subsection 3.3) to better understand its interplay with potential Y1 outcomes.

4 Additional Recommendations

A small change to the southern portion of the footprint improves overlap with the Euclid footprint and determines negligible changes in science metrics.

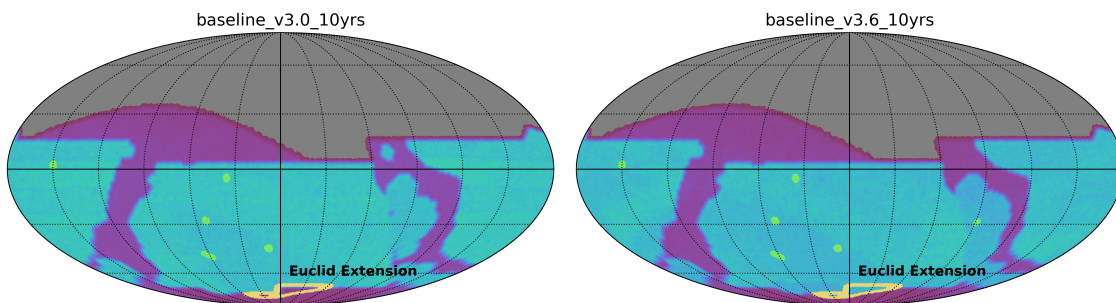


FIGURE 13: Small changes to the southern portion of the footprint improve overlap with Euclid.

The airmass limits for the Near-Sun Twilight microsurvey, introduced with baseline v3.0, were increased from $X = 2.5$ to $X = 3.0$ in v3.2, corresponding with decreasing the minimum solar elongation reached for this microsurvey from 60 degrees to 45 degrees. This improves the likelihood of discovery of interior-to-earth objects, increasing the survey sensitivity to this niche of discovery space. The recovered population of objects interior to Venus at magnitude $H \leq 20$ goes from $\sim 4\%$ to $\sim 40\%$ in v3.2 and later. The impacts outside the microsurvey are negligible.

5 Additional changes introduced in v3.6 OpSims

Starting with v3.6 some important assumptions underlying the simulations were updated:

- Increased downtime in Y1 to reflect a more realistic transition into operations. The downtime in Y1 is simulated to be maximal early on and decreased to the level expected for the general LSST survey by the end of the first year.
- The effect of jerk on slew time is now included in the simulations, and thus included in scheduling choices.

6 Summary of SCOC Phase 3 recommendations

The following list summarizes the Phase 3 recommendations for the LSST survey strategy which together with the recommendations in PSTN-053 and PSTN-055 define the LSST survey strategy starting plan. The recommendations included in this report are listed below.

- The SCOC recommends swapping u - and y -band according to the moon phase. This produces benefits for SN cosmology while preserving coverage on short timescales. This recommendation is implemented starting in `baseline_v3.2` (subsection 3.1).
- Following updates to the mirror coating plans, with all three mirrors in the system to be coated in silver, which increases throughput in all bands bluer than u but decreases u throughput (by $\lesssim 30\%$), the SCOC recommends an increase of the exposure time in u -band to 38 seconds per visit and an increase of the number of u -band visits of 10% compared to `baseline_v3.0`. To compensate for the excess time dedicated to u band the

SCOC recommends decreasing the exposure time in the other bands. Simulations show that this comports a decrease of 0.8 second exposure time in all other bands (subsection 3.2).

- The SCOC iterates the recommendation for rolling in two sky areas at 0.9 strength on the WFD low-dust footprint. We are adopting the rolling uniform strategy designed by the Uniformity Task Force in `baseline_v4.0` simulations, which implements three cycles of rolling, but we will continue to investigate implementations of rolling until our Y1 recommendation since rolling will under no implementation under consideration start before Y2. The SCOC recommends that the time domain community, particularly those interested in phenomena that have evolutionary time scales of hours-to-days, urgently quantifies the impact of the proposed uniform rolling compared to rolling in four cycles. Further, the SCOC iterates the recommendation that Data Management scopes a plan for producing uniform data releases in DR5 and DR8 in addition to the standard data releases and that the cost of the development and storage of these additional data releases is scoped and shared with the scientific community (subsection 3.3).
- The SCOC concluded that rolling on the Galactic footprint would have a net negative effect on the survey as a whole and would benefit Galactic science, and recommends no rolling in the Plane or Bulge (subsubsection 3.4.1).
- The SCOC recommends redistributing the excess visits in the “blob” centered around a Galactic longitude of $l = +45$ to cover a low-visit “barrier” at $l = +335$ in the Plane and at the border of the Plane and Bulge. This change would give continuous longitude coverage along the Plane from a longitude of $l = +30$ down through $l = +280$ and boost metrics for time-domain science in the Bulge/Plane (subsubsection 3.4.1).
- The SCOC recommends that a visit plan consistent with this `roman_v3.3` with the number of redistributed visits to be capped at that used in this simulation: in the current Roman field `OpSims` the number of visit reallocated to the Roman field is $\sim 1,600$. However, the timing of the implementation of this augmented observing campaign needs to remain flexible at this time to respond to the as-of-yet to be finalized launch date of Roman and the scheduling of its surveys (subsubsection 3.4.1).
- The SCOC finds that the adoption of a revised filter balance in the Bulge and Plane with less y and more z , g , and u compared to the present baseline is potentially beneficial on the net, but that existing metrics are not adequately sensitive to the explored filter balance changes for some expected science cases. The SCOC concludes that a survey using the currently implemented filter balance in the Bulge and Plane in `baseline_v3.4`

will produce excellent science and the LSST can start with this implementation (subsubsection 3.4.2).

- The SCOC recommends a bluer filter mix in the SMC, LMC, and SCP regions, bounded by the requirement that the increased number of dark-time visits in a relatively narrow range of right ascension does not affect other parts of the survey (subsubsection 3.4.3).
- The SCOC recommends the implementation of a LSST ToO program as detailed in *Rubin ToO 2024: Envisioning the Vera C. Rubin Observatory LSST Target of Opportunity program*¹⁵ by the scientific community at large (subsection 3.5).
- The SCOC recommends that a meeting to follow Rubin ToO 2024 is organized closer to the start of O5 to refine the GW follow-up survey strategy with improved knowledge of the expected performance of the GW detector networks and systems in O5 and of the performance of the full Rubin system and that the solar system and neutrino ToOs should start as soon as possible: as soon as suitable templates are available (subsection 3.5).
- The SCOC recommends that Rubin only consider potential ToOs that emanate from vetted discovery and distribution systems that produce and dispatch fully machine-readable alerts (subsection 3.5).
- The SCOC recommends real-time human review of potential ToO triggers and the establishment of a Rubin ToO Advisory Committee as described above (subsection 3.5).
- The SCOC recommends that, if the technical feasibility is confirmed in commissioning, the survey is conducted with single exposures. With our recommendation of modifying the exposure time for u band to 38 seconds, and compensating for this extra u band survey time by short decrease in exposure across all other bands, the single visits would be 1x29.2 seconds (subsection 3.6).
- The SCOC recommends that DDF observations should be sequences of multiple WFD-like visits (as opposed to increased exposure times) to allow rapid alert generation (subsection 3.7).
- The SCOC recommends that the baseline translational dithering scale of DDF observations should be reduced from 0.7 degrees to 0.2 degrees (with exploration of even smaller translational dithers compatible with instrumental signature removal and calibration needs) (subsection 3.7).

¹⁵https://docs.google.com/document/d/1WE4NG13dF0VGo7lZpyG1fe_JiX9m-kL15JYQkhu9iso/edit?usp=sharing

- The SCOC recommends that the baseline survey strategy should accommodate varying the nightly depth, filters, or cadence of different DDFs throughout the course of LSST, while maintaining the Phase 2 recommendations for the 10-year depth of each field (including the enhanced COSMOS observations to reach 10-year depth in the first 3 years) (subsection 3.7).
- The SCOC urges the Data Management and Alert Production teams to assess the feasibility of, and resources needed for, enabling nightly co-adds of sequential DDF visits and recommends that a path is developed to enable the creation of these co-adds, subtraction with deep templates, and faint alert generation (with higher latency as needed, e.g., after sunrise) (subsection 3.7).
- The SCOC recommends the airmass limit for the Near-Sun Twilight microsurvey is increased to $X = 3.0$ (section 4).
- The SCOC recommends a slight modification of the `baseline_v3.0` footprint to improve overlap with the Euclid footprint (section 4).

These recommendations will be implemented in the `baseline_v4.0` simulations. A set of simulations tagged `v3.6` is made available for the community to assess the impact of different aspects of the recommendation. Note that all of these simulations include the updated downtime and effects of slew jerk.

Figure 14 shows the performance of the survey strategy on a set of core LSST science and system metrics. Note: Significant improvements were obtained on most metrics through `v3.0`. Those are to be attributed to changes of the survey strategy through community input and SCOC recommendations. The visible improvement on nearly all metrics between `baseline_v3.2` and `baseline_v3.3` is attributed to the updated filter transmission curves. The survey strategy is largely unchanged between `baseline_v3.3` and `baseline_v3.4`; the small changes in performance are to be attributed to `rubin_scheduler` code updates.¹⁶ `baseline_v3.6` reflects this recommendations. The overall apparent drop in performance between `baseline_v3.4` and `baseline_v3.6` is primarily due to the inclusion of slew time jerk effects and more realistic estimates of downtime in Y1 (section 5). We also make available a variation of `baseline_v3.6` with four cycles of rolling to enable the investigations of different rolling implementations. While our recommendation is to implement the ToO program as described in subsection 3.5, we provide an `OpSim` consistent with `baseline_v3.6`

¹⁶See <https://survey-strategy.lsst.io/baseline/changes.html>.

but without the ToO program to allow the community to see the small effects that the introduction of ToOs has on the LSST. Finally we provide an implementation of `baseline_v3.6` with single exposure visits (instead of 2x15 second snaps, subsection 3.6) which, pending commissioning outcomes, is the expected observing mode. In this `OpSim`, the survey time gained by dropping snaps (decreased readtime per visit) is allocated evenly across all all-sky observing modes: this includes the WFD, NES, SCP and Galactic Plane.

Draft




FIGURE 14: LSST performance on key metrics for different OpSims including the Phase 3 recommendation (this document). From left to right, then sequence shows progressively newer OpSims, starting with the baseline in effect through 2018 (v1.x) up to the baseline proposed with this recommendation (baseline_v3.6) and some variations on the latter. Starting with baseline_v3.6 we include more realistic down time expectations in Y1 and the effects of slew jerk on scheduling and observing. This causes an overall decrease in the number of visits (see also Figure 15). baseline_v3.6 is also the first OpSim shown here which includes the ToO program. Rolling is implemented in 3 cycles in baseline_v3.6. This achieves desired uniformity in DR5 and DR8 (see subsection 3.3). four_cycles implements rolling in four cycles. The SCOC will not commit to a recommendation on the specific rolling implementation until the release of its Y1 recommendation. The community is encouraged to explore the impact of adopting either rolling strategy on their science. A version of baseline_v3.6 without ToO (no ToO) is included to allow the community to see the impact of the ToO program but the SCOC is committed to recommending the implementation of a ToO program as simulated in baseline_v3.6. Finally, we include an implementation of baseline_v3.6 where visits are conducted in a single exposure instead of two snaps (single_snap). The SCOC recommends the implementation of LSST in single visits (as shown in this simulation); this recommendation is, however, pending commissioning outcomes.




FIGURE 15: Number of LSST visits for the same set of `OpSims` used in Figure 14. The number of observations decreases between `v3.4` and `baseline_v3.6` because of the inclusion of more realistic down times in Y1, of the effects of jerk on slew and scheduling, and, in part, because of the inclusion of the ToO program in LSST. This drop in number of visit determines a general loss of performance on all metrics, as seen in Figure 14. The increase in efficiency associated with moving to single exposure visits (the `OpSim` labeled as `single snaps`) largely recovers the visits lost between `3.4` and `baseline_v3.6`. Nonetheless, after Y1, when the system performance is better understood, in the future, the SCOC will consider how the additional time may be allocated (subsection 3.6).

A References

- Abbott, T.M.C., Aguena, M., Alarcon, A., et al., 2022, *Phys. Rev. D*, 105, 023520 (arXiv:2105.13549), doi:10.1103/PhysRevD.105.023520, ADS Link
- Baleato Lizancos, A., White, M., 2023, *Journal of Cosmology and Astroparticle Physics*, 2023, 044 (arXiv:2305.15406), doi:10.1088/1475-7516/2023/07/044, ADS Link
- Brandt, W.N., Ni, Q., Yang, G., et al., 2018, arXiv e-prints, arXiv:1811.06542 (arXiv:1811.06542), doi:10.48550/arXiv.1811.06542, ADS Link
- Coughlin, M.W., Dietrich, T., Heinzl, J., et al., 2020, *Phys. Rev. Res.*, 2, 022006, URL <https://link.aps.org/doi/10.1103/PhysRevResearch.2.022006>, doi:10.1103/PhysRevResearch.2.022006
- [PSTN-055]**, F. Bianco and The Rubin Observatory Survey Cadence Optimization Committee, 2023, Survey Cadence Optimization Committee's Phase 2 Recommendations, URL <https://>

//pstn-055.lsst.io/,

Vera C. Rubin Observatory Project Science Technical Note PSTN-055

Gianfagna, G., Piro, L., Pannarale, F., et al., 2024, Monthly Notices of the Royal Astronomical Society, 528, 2600

Graham, M.L., Connolly, A.J., Ivezić, Z., et al., 2017, The Astronomical Journal, 155, 1, URL <http://dx.doi.org/10.3847/1538-3881/aa99d4>, doi:10.3847/1538-3881/aa99d4

[LPM-17], Ivezić, Ž., The LSST Science Collaboration, 2018, LSST Science Requirements Document, URL <https://ls.st/LPM-17>, Vera C. Rubin Observatory LPM-17

Lochner, M., Scolnic, D., Almoubayyed, H., et al., 2022, ApJS, 259, 58 (arXiv:2104.05676), doi:10.3847/1538-4365/ac5033, ADS Link

LSST Science Collaboration, 2009, ArXiv e-prints (arXiv:0912.0201), doi:10.48550/arXiv.0912.0201, ADS Link

Scolnic, D.M., Lochner, M., Gris, P., et al., 2018, arXiv e-prints, arXiv:1812.00516 (arXiv:1812.00516), doi:10.48550/arXiv.1812.00516, ADS Link

[PSTN-053], Z. Ivezić and The Rubin Observatory Survey Cadence Optimization Committee, 2022, Survey Cadence Optimization Committee's Phase 1 Recommendation, URL <https://pstn-053.lsst.io/>, Vera C. Rubin Observatory Project Science Technical Note PSTN-053

B Acronyms

Acronym	Description
AGN	Active Galactic Nuclei
B	Byte (8 bit)
DDF	Deep Drilling Field
DESC	Dark Energy Science Collaboration
DM	Data Management
DR11	Data Release 11
DR2	Data Release 2
GW	Gravitational Wave

JPL	Jet Propulsion Laboratory (DE ephemerides)
LIGO	Laser Interferometer Gravitational-Wave Observatory
LMC	Large Magellanic Cloud
LPM	LSST Project Management (Document Handle)
LSST	Legacy Survey of Space and Time (formerly Large Synoptic Survey Telescope)
M1	primary mirror
M2	Secondary Mirror
M3	tertiary mirror
MAF	Metric Analysis Framework
MC	Monte-Carlo (simulation/process)
MMA	Multi Messenger Astronomy
MW	Milky Way
NES	North Ecliptic Spur
PCW	Project Community Workshop
PST	Project Science Team
PSTN	Project Science Technical Note
RA	Risk Assessment
SC	Science Collaboration
SCOC	Survey Cadence Optimization Committee
SCP	South Celestial Pool
SMC	Small Magellanic Cloud
SN	SuperNovae
SRD	LSST Science Requirements; LPM-17
TVS	Transients and Variable Stars Science Collaboration
ToO	Target of Opportunity
WFD	Wide Fast Deep
YSO	Young Stellar Object
photo-z	photometric redshift