LSST Cadence Notes Solicitation

Survey Cadence Optimization Committee and the Survey Scheduling Team
December 8, 2020
Updated Feb 11, 2021

This solicitation calls for input (submission details are listed at the end of this document) to further evaluate and explore survey strategy options presented in the Fall 2020 LSST Cadence Optimization Report, “Survey Strategy and Cadence Choices for the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST)” (PSTN-051, https://pstn-051.lsst.io/). The report compiles a large series of simulations exploring various survey strategy options, in response to white papers submitted as part of the 2018 Call for White Papers on the LSST Cadence Optimization (https://www.lsst.org/call-whitepaper-2018) and the advisory report from the LSST SAC (https://ls.st/document-32816). However, the report could only evaluate the available simulations using metrics contributed by that time. *We would like to take this opportunity to fill additional gaps in metric coverage and get feedback on the existing large set of simulations, based on responses to six specific questions.*

As stated in the 2018 Call for White Papers, the LSST baseline cadence is an existence proof that LSST dataset can be delivered as designed and advertised (for more details, see sections 2.1.5, 2.2.2 and 3.1 in the LSST overview paper, ls.st/lop). The simulations explored in the Fall 2020 Cadence Report offered variations on this survey strategy in an attempt to further enhance the baseline science yield from LSST. It is noteworthy that this optimization is essentially fine tuning, rather than starting from scratch - anticipated gains in various metrics are closer to 10% than, e.g., a factor of two.

What we expect, and what we see in the metrics run in the standard MAF analysis, is that science enhancements can be accomplished by varying at least some of the fundamental survey parameters, including:

(1) Survey footprint and distribution of visits
(2) Exposure time per visit
(3) Allocation of observing time per band - the distribution of visits between filters
(4) Time sampling (cadence) and dithers - on timescales from nightly to monthly to yearly

In each case, we first briefly summarize the rationale behind the baseline cadence parameters and then formulate a list of specific questions to be responded to in submitted Cadence Notes. The reasoning behind these questions, and behind accompanying assertions, is based on the simulations included in the Fall 2020 Cadence Report.

High level summaries of the simulations, including links to download the databases and links to the MAF outputs, are available at
https://community.lsst.org/t/fbs-1-5-release-may-update-bonus-fbs-1-5-release ,
https://community.lsst.org/t/fbs-1-6-release-august-2020 **and**
https://community.lsst.org/t/survey-simulations-v1-7-release-january-2021/

For an executive summary, please see this document.

The principal parameters that define an observing strategy include:

A) Survey footprint: How the survey visits are distributed across the sky. The footprint consists of one or more regions on the sky that share a common observing pattern.

The baseline survey footprint includes the main Wide-Fast-Deep (WFD) survey, several mini-surveys and the so-called Deep Drilling fields.

WFD footprint definition: The baseline survey footprint uses simple cuts in galactic longitude and latitude, together with Declination cuts, to define the 18,000 sq. deg. of the WFD. Instead, we suggest moving to dust extinction limits to define borders rather than simple cuts in galactic longitude and latitude, and varying the Declination limits to adjust the overall area. Using galaxy counts as a proxy for extragalactic science, we will eventually optimize the dust extinction threshold to maximize the number of galaxies above some SNR cut. Note that apparently we are already very close to the optimum. Nevertheless, because the counts vs. area curve has a very broad maximum, we still need to decide whether to push beyond the SRD (LSST Science Requirements Document, ls.st/srd) design requirement of 18,000 sq. deg. and perhaps extend the WFD area to as much as 20,000 sq.deg. (with a corresponding decrease of the number of visits per pointing).

WFD observing time allocation: The SRD assumes about 90% of the total survey time would be spent on the WFD; however, it is plausible that 75%-80% might suffice to meet the SRD requirements - *if the current system performance estimates hold up*. The four approaches to utilize this hypothetically additional observing time include:

- a) increase the number of visits for the main WFD survey,
- b) increase the area for the main WFD survey (or combine with a)
- c) use it for mini-surveys, and
- d) use it for the so-called Deep Drilling Fields.

Mini-surveys, defined as regions larger than a single field, with footprint or cadence that differ from the main WFD survey, currently include: the Galactic plane (GP) mini-survey, the North Ecliptic Survey (NES), and the South Galactic Pole (SCP). The boundaries for the NES and the SCP are relatively well defined (once the WFD region is defined), while the GP mini-survey has multiple further options. *The cadence and number of visits available for these mini-surveys are some of the key open questions.*
As an added complication, optimization of crowded fields in the GP will have to be coupled with the performance estimates for DM science pipelines.

Deep Drilling Fields: given a fraction of observing time dedicated to Deep Drilling Fields (DDFs), details of DDFs strategies don't have a strong influence on the rest of the survey. The current DDF implementation generally reflects the science cases proposed in the last call for

DDF white papers. The amount of time devoted to DDFs is generally held steady at 5%; this allows 1% of the survey time to be spent on each of the five DDFs. The location of the first four DDFs has been fixed for several years, to match existing multi-wavelength deep survey locations. There is an informal agreement with the Euclid survey for the fifth DDF field to overlap the Euclid Deep Field-South.

Q1: Are there any science drivers that would strongly argue for, or against, increasing the WFD footprint from 18,000 sq. deg. to 20,000 sq.deg.? Note that the resulting number of visits per pointing would drop by about 10%. If available, please mention specific simulated cadences, and specific metrics, that support your answer.

Q2: Assuming that current system performance estimates will hold up, we plan to utilize the additional observing time (which may be as much as 10% of the survey observing time) for visits for the mini-surveys and the DDFs (with an implicit assumption that the main WFD survey meeting SRD requirements will always be the first priority). What is the best scientific use of this time? If available, please mention specific simulated cadences, and specific metrics, that support your answer.

Relevant simulations for Q1 and Q2 would include the **wfd_scale** family (where the area of the WFD is held constant, but the number of visits per pointing - equivalent to the weight on the WFD - is varied), the **footprints** family (where the overall survey footprint, including the location of the WFD is varied), the **filter_dist** family (where the area of the WFD is increased, at the cost of all other mini-surveys) and the v1.7 family **footprint_tune**.

B) Exposure time per visit

Given the system characteristics, and assuming the same exposure time for all bands and all visits, the optimal exposure time per visit is in the range 20 sec to 40 sec (for science-driven arguments, see Section 2.2.2. in the LSST overview paper). The baseline cadence assumes 30 sec per visit, split into two readouts (2x15 sec) to enable efficient detection of cosmic rays and rapid variability (by subtracting the two 15 sec exposures). While a single readout of 30 sec visit exposure would be more efficient
(providing about 7% more open-shutter time and an overall gain of 9% in number of visits), for programmatic reasons we will retain 2x15 sec visits at least into the commissioning phase.

One possible exception is the u band, where the combination of camera readout noise and the low sky background argues for longer exposures. We are currently considering changing the u band exposure time to a single 50 sec readout, which would result in about 0.5 mag deeper limiting magnitudes (both for single visits and for co-added data). Such significantly deeper u band data would have a number of scientific benefits, most notably improved photometric redshift and photometric metallicity performance, as well as improved transient classification. To maintain the frequency and overall number of u band visits, the extra 5% of observing time (i.e. the change from 7% to 12% allocated to the u band) would be shifted from other bands, resulting in up to about 0.1 mag shallower coadded images, and 5%-10% fewer visits in the grizy bands.

Q3: Are there any science drivers that would strongly argue for, or against, the proposal to change the u band exposure from 2x15 sec to 1x50 sec? If available, please mention specific simulated cadences, and specific metrics, that support your answer.

Relevant simulations for Q3 would include the **baseline** family and the **u_long** family.

C) Allocation of observing time per band

The SRD suggests for the main WFD survey in Table 24 the following fractions of observing time per band: 0.07, 0.10, 0.22, 0.22, 0.19, 0.19 in ugrizy, respectively.
The change of u band exposure from 2x15 sec to 1x50 sec would modify this allocation to 0.12, 0.09, 0.21, 0.21, 0.18, 0.18.

Q4: Are there any science drivers that would strongly argue for, or against, further changes in observing time allocation per band (e.g., skewed much more towards the blue or the red side of the spectrum)? If available, please mention specific simulated cadences, and specific metrics, that support your answer.

The most televant simulations for Q4 are the **filter_dist** family.

D) Time sampling (cadence) and revisit offsets (translational and rotational dithers)

The cadence for observing each pointing has options on ranges from within each night, to across a season (the season length), and then from year to year (rolling cadence).

The baseline cadence attempts to acquire two visits per night, separated by about 30 minutes. Given a pair of visits, *an open question is whether to use the same filter, or different filters.* As yet another variation, three visits per night for some fraction of the fields observed per night, with two in the same filter to measure brightness change, and the third in a different filter to measure transient's color, remains a viable strategy.

From night to night, it is possible to limit visits to sub-regions of the sky. The **AltSched** scheduling algorithm is an example of this; on alternating nights, the AltSched algorithm chooses either a north or south region of the sky to observe; this can reinforce particular revisit times (every other night) while reducing others. Within a month, we naturally tend to use bluer filters closer to new moon and redder filters near the full moon; we can emphasize this to a greater or lesser extent. Over a season, if visits are distributed uniformly throughout the observing season, the season length and typical revisit time are strongly correlated; longer seasons and higher hour angle limits correspond to longer revisit times. Across years, the

rolling cadence sampling strategy enhances the number of visits for a chosen area of sky in a given year resulting in a decreased revisit time, at the expense of a longer revisit time in other years. This results in a net effect of the loss of transients on time scales comparable to a field being ``inactive'', but during active periods with denser temporal sampling produces better sampled light curves. This strategy was driven by supernovae science and accomplishes its goals, but the coverage gaps may be too long for variable sources such as quasars.

Revisit offsets -- dithers -- can be considered in terms of both translational and rotational dither patterns, and can be implemented differently in the all-sky portions of the survey versus focused fields like the DDFs. The current baseline for translational dithers over the general sky is to maintain a consistent tiling within each night, then randomly rotate the tiling between nights; this allows the pairs (or triplets) of visits to contain the same sky area. For DDFs, we implement fairly large translational dithers to aid calibration.

Q5: Are there any science drivers that would strongly argue for, or against, obtained two visits in a pair in the same (or different) filter? Or the benefits or drawbacks of dedicating a portion of each night to obtaining a third (triplet) visit? If available, please mention specific simulated cadences, and specific metrics, that support your answer.

Q6: Are there any science drivers that would strongly argue for, or against, the rolling cadence scenario? Or for or against varying the season length? Or for or against the AltSched N/S nightly pattern of visits? If available, please mention specific simulated cadences, and specific metrics, that support your answer.

Q7: Are there any science drivers pushing for or against particular dithering patterns (either rotational dithers or translational dithers?) If available, please mention specific simulated cadences, and specific metrics, that support your answer.

Relevant simulations for Q5: the nightly cadence is varied in the **baseline** family (baseline and baseline_samefilt), while a triplet visit is added in the **third_obs** family. For Q6: various rolling cadences are simulated in the **rolling** family (v1.7), and AltSched pattern visits are used in the **alt_rolling** family. Varying the timing of visits in different bandpasses throughout the month is explored in the **filter_cadence** family (v1.7). A range of potential DDF dither options are simulated in the **ddf_dither** and **euclid_dither** families.

* The families referenced above are intended to be helpful suggestions: you do not have to constrain your answers to those simulations only.

Answers to these questions can be of arbitrary length, but we anticipate that about 3 pages ought to be sufficient. Please try to address our questions within the suggested 3 page limit, but please don't skip meaningful points just so that your Cadence Note doesn't exceed this limit. We will **not** reject any submissions for violating the page limit. If it seems that what you have to say would take much longer than 3 pages, perhaps you need to submit more than one Cadence

Note. If unsure, don’t agonize over it but describe your case to Zeljko in an email (ivezic@uw.edu) and he’ll discuss it with the SCOC and get back to you promptly.

Cadence Notes can be submitted on behalf of Science Collaborations, or by individuals. We are hopeful that all Science Collaborations will submit their Note, even in cases when the existing baseline cadence fully meets science needs. According to our current plan, Cadence Notes are the last formal document solicitation. Nevertheless, the SCOC wants to be inclusive and additional unsolicited input will be accepted, including discussions of the SCOC liaisons with their Science Collaborations.

Please send pdf documents by midnight, Pacific time zone, **April 15, 2021** (updated from Mar 15) to ivezic@uw.edu with the Subject: line “LSST Cadence Note”.