A Report from the LSST Science Advisory Committee: Recommendations for Operations Simulator Experiments Based on Submitted Cadence Optimization White Papers

The LSST Science Advisory Committee

April 2019

1 Introduction

With its large telescope aperture and field of view, the Large Synoptic Survey Telescope (LSST) will be the most powerful visible-light survey facility in the history of astronomy. The large aperture (effective diameter of 6.7 meters) means that it can go faint with very short exposure times. The wide field of view of the imaging camera (9.6 deg^2) means that it can quickly cover wide swathes of the sky. The fact that the telescope will be in continuous operation for ten years means that it can do so repeatedly. This Wide-Fast-Deep (WFD) capability enables an enormous range of scientific opportunities, as has been detailed by the LSST overview paper (Ivezić et al. 2019) and the LSST Science Book. Given this combination of wide field of view and large aperture, much of this science can be done with a single observing plan. Indeed, the LSST Science Requirements Document (SRD; Ivezić et al.) lists four science themes that LSST is designed to address:

1. Constraining Dark Energy and Dark Matter.
2. Taking an Inventory of the Solar System.
3. Exploring the Transient Optical Sky.
4. Mapping the Milky Way.

In preparation for the LSST Final Design Review before construction was approved, the LSST Project developed a so-called "universal cadence" plan, in which the sky is tiled in a given filter to maximize sky coverage over the WFD footprint (defined by cuts in declination and Galactic latitude) before changing to the next filter. With this strategy, each field in the footprint will be observed over 800 times, summed across the six filters, over the ten-year span of the survey. Additional survey areas (at low Galactic latitude, low ecliptic latitudes in the North, the South Celestial Pole) have observations at reduced cadence, and four so-called Deep Drilling Fields (DDFs) are single LSST pointings that receive considerably more exposures and thus go deeper than the main WFD footprint. A full description of the universal cadence may be found in the Community Observing Strategy Evaluation Paper (LSST Science Collaborations et al. 2017, hereafter COSEP). With this universal cadence, the Project was able to demonstrate that the LSST indeed could address the four science themes. But it was clear that the universal cadence was not optimal, in the sense of maximizing the science opportunities both in the four science themes above, and the myriad of other areas in which LSST is likely to have an impact.

To address this concern, the LSST science community was asked to prepare white papers with suggestions for alternatives to the baseline cadence. The 46 white papers that were received advocate for a variety of scientific opportunities that LSST enables, and specific observing strategies, or cadences, that will support that science. Given the ten-year span of the LSST survey, there are inevitable tradeoffs. For example, while the universal cadence does a good job of covering the sky relatively uniformly with time, it is not ideal from the point of view of measuring variable phenomena on the timescale of hours to days. While those interested in extragalactic science will want to maximize the telescope time spent at high Galactic latitudes, there is a great deal of both static and time-variable science to be explored with extensive observations in the Galactic plane. Understanding the tradeoffs, both practical and scientific, requires extensive simulations of the LSST observing strategy under different rules governing how the sky is tiled. The LSST Science Advisory Committee was charged with recommending simulation experiments to be done based on these white papers. These simulations will be the raw material for making decisions on the algorithms that drive the LSST observing strategy. These recommendations are going to the LSST Operations Simulator (OpSim) team, including Lynne Jones, Peter Yoachim, and Tiago Ribeiro. We see this as an iterative process; as OpSim outputs are made available to the LSST science community, the results will inform decisions about refined simulation experiments. Thus we anticipate updating this document as we learn more about the performance of the system. We encourage members of the science community, especially the LSST Science Collaborations, to give feedback on the OpSim results and to suggest further experiments. As OpSim runs are produced, they will be announced on `https://community.lsst.org/c/sci/survey-strategy`, which is a good forum for discussion. As this discussion converges on specific recommendations from the various LSST science communities, please post summaries via a google form at `https://tinyurl.com/OpSim-feedback`.

In § 2, we briefly describe the call for white papers from the community, and the process by which the SAC reviewed and organized its recommendations. We make high-level comments in § 3. Our detailed recommendations may be found in § 4, and we conclude in § 5. We include an Appendix listing the 46 white papers that were received.

2 The Process

In June 2018 a call was made to the world-wide community of scientists (Ivezić et al. 2018) for white papers giving science cases that LSST could address and suggesting specific observing strategies to optimize that science. The scope of these strategies could be as narrow as a specific pointing in which deep observations could be made with a specific cadence, or as broad as a re-definition of the footprint of the LSST WFD survey and the strategy with which it is observed. White paper authors were asked to respect the scientific requirements of the project, as outlined in the SRD. These white papers were due on November 30, 2018; 46 such papers were received from all over the world. The authors were dominated by, but by no means restricted to, the membership of the LSST Science Collaborations. The white papers at least touched upon most, if not all, major LSST science themes; the SAC did not identify any major topics that had been discussed in the community that were not addressed by at least one white paper.

We found it useful to sort the white papers into a set of common themes, to organize our review and discussion. They were as follows:

- Suggestions for the footprint of the main Wide-Fast-Deep survey. These principally examined the definition of the Northern and Southern boundaries of the footprint, as well as the division, if any, between high-latitude and low-latitude sky.
- Specific cadence suggestions for the WFD survey. Among other things, this included considerations of the nature of individual visits (the current default is that these would each consist of two 15-second exposures), the choice of filters as a function of time, and the tradeoff between covering the entire available sky uniformly each year of the survey, and concentrating on limited areas each year to increase the cadence.
- Cadence suggestions for observations at low Galactic latitude.
- Suggestions for "mini-surveys", regions of sky that would be observed with a different cadence and/or depth from the WFD survey.
- Suggestions for new deep-drilling fields, and cadences with which both new and old ones would be observed.
- Specialized observing modes, including twilight observing, target of opportunity observations, very short exposures, and untracked observations.

We include the full list of white papers in the Appendix, ordered roughly as described above, reflecting the order in which we carried out our review.

The SAC consists of 15 members. Each white paper was assigned to four SAC members to read and comment on (thus each SAC member closely read of order a dozen white papers), in preparation for a two-day detailed face-to-face review of the whole set of white papers. The white papers were grouped into common themes that were discussed together during that two-day meeting. The collected notes from the meeting, and subsequent phone conferences, were the raw material from which the current report was prepared. We also reached out to a number of white paper authors and science collaborations for clarification on some aspects of their white papers.

As emphasized above, as results from the OpSim experiments we suggest become available, the scientific community will have ideas for refinements and further experiments to be carried out. The full set of experiments will have been run by December 2019. The LSST operations project will then stand up a Survey Strategy Committee, which will be charged with using the results of these experiments to make recommendations for the survey strategy to be used at the start of full LSST operations.

3 General Considerations

The LSST Operations Simulator (OpSim) is described in Chapter 2 of the COSEP (also see Delgado et al. 2014). It includes a detailed model for the motion of the telescope (allowing slew times from one field to another to be simulated accurately), overheads associated with camera reads and filter changes, the effects of twilight and the moon, a model for the weather, and myriad other factors that affect observing efficiency. Given various rules governing the sky coverage of LSST, it can simulate the roughly 1.8 million visits that LSST will carry out over its 10-year lifetime, together with detailed statistics describing, e.g., the coadded depth for each filter as a function of position across the survey, the cadence with which those observations will be carried out, and so on. The OpSim code has been extensively rewritten since the publication of the COSEP: it is now significantly more flexible and able to simulate far more complex optimizations, and can simulate a full 10-year survey in less than 7 hours (rather than the several days previously required), allowing many more simulations to be run. This new version of the OpSim code contains the prototype for the LSST Scheduler, which will run during the survey itself, and determine on a nightly, monthly and yearly basis what the pointing sequence will be.

As the SAC understands it, the current version of the OpSim code is sufficiently powerful to carry out the majority of the experiments we suggest in § 4. However, there are two capabilities that will need to be developed going forward:

- Scheduling specific cadences tied to other facilities. Several white papers described observing modes in which the LSST coverage of a given area of sky was timed to coincide with its coverage with another facility, such as WFIRST or Euclid. In practice, this will require giving OpSim (and ultimately, the Scheduler software) specific dates or times to observe specific fields. This capability needs to be built into OpSim.
- OpSim currently uses only one 10-year realization of the weather. This realization is realistic, as it is based on detailed weather records from Cerro Tololo Observatory ($<$ 20 kilometers from Cerro Pachón, the peak on which LSST sits). However, we need to understand the variance of OpSim outputs due to different weather patterns. For example, how would LSST science be affected if the first year had particularly

bad weather? The SAC recommends that there be a way to randomize the assumed weather model over the ten years (perhaps by scrambling the years of the current weather realization).

Once OpSim is run, one needs to assess how well a given realization of the 10-year survey addresses the many science goals that we want to keep in mind. Some of these goals are quantified in the SRD, while many of those highlighted in the white papers explore other aspects of LSST science. We need to quantify the performance of each OpSim run in terms of these goals. This is done via the Metrics Analysis Framework (MAF), as described in the COSEP (see also Jones et al. 2014). The MAF consists of algorithms run on each OpSim realization of the 10-year survey that measure the performance, via quantitative metrics, of that realization for a given science goal or requirement. This allows a direct determination of the scientific impact of different cadence algorithms and assumptions. MAF algorithms already exist for essentially all the requirements listed in the SRD, as well as some of the science goals listed in the cadence optimization white papers. But there is quite a bit more work to be done: many of the white papers only outline what an appropriate metric would be, and coding these up will require substantial effort. The SAC is concerned about how this work will be done; experience has shown that metrics are often only coded up when the OpSim team is closely involved in the effort. We recommend that the OpSim team be given the resources to code the metrics suggested in the white papers; the SAC is happy to work with them to identify those metrics we find most urgent.

4 Our Recommendations

Our goal here is to produce a list of experiments to be run with the Operations Simulator, that will nominally satisfy the core science goals of LSST. Our inspiration comes from the cadence optimization white papers. While many of the specifics can be traced to individual white papers (which we quote explicitly on occasion), we have carried out a synthesis of the ideas in the white papers in many cases.

At the level that the SAC's deliberations were carried out, a specific OpSim run is specified by:

- The assumed footprint of the Wide-Fast-Deep survey;
- The cadence of the WFD observations, including any rolling of the sky coverage.
- Mini-surveys (extended regions larger than a Deep Drilling field, to be observed in a different mode or cadence from WFD). The definition of these will depend in part on the definition of WFD, as many of the suggested mini-surveys extend the footprint of WFD (further North, further South, or to lower Galactic latitudes).
- Deep Drilling fields, and their cadence.
- Any further observing programs, such as Target of Opportunity.

With this in mind, we will discuss the specific options we have in each of these categories. The large number of options in each category leads to a very large number of possible combinations. Not all combinations will lead to surveys that respect the core LSST science goals, or that can be done in the ten years that LSST have available. The OpSim team has developed tools to quickly estimate how feasible a given combination is in this sense; it may be that some of our suggested combinations will not be workable. Furthermore, there are many specific questions, e.g., about the effects of certain cadence decisions or choices of Deep Drilling Fields, that can be addressed with focused OpSim experiments on an otherwise fiducial OpSim run.

The real limitation in practice is the time needed to understand and interpret the outputs of each simulation. Metrics capturing all the different science goals in the SRD and the white papers need to be compared, to determine what the observing efficiency and scientific tradeoffs really are.

4.1 The WFD Footprint

The LSST default definition of the WFD footprint may be found in the LSST Science Book (Chapter 3) and the LSST overview paper. It is characterized by strict limits in declination: $-62^\circ < \delta < +2^\circ$, to minimize observations at high airmass. It also has a cut at low Galactic latitudes, which is designed to exclude regions of particularly high stellar density; this is driven in part by concerns about the performance of the LSST image-processing pipeline when the stellar density becomes too high.

There were a number of white papers that argued for variations of this footprint. The appropriate cut at low Galactic latitudes depends on one's goals: a cut based on stellar density makes sense if the concern is image processing (as above) or avoiding the confusion limit in coadded depth. A cut based on Milky Way extinction makes sense if one's goal is to maximize the number of galaxies in the WFD footprint. No Galactic latitude cut at all makes sense if the goal is to maximize the number of stars in the survey.

Similarly, there were cogent arguments to extend the WFD footprint further North and further South. Based on these arguments, we suggest the following options for the WFD footprint in the OpSim experiments:

1. The original footprint, as defined in the current baseline OpSim runs. Again, it extends in declination from $\delta = -62^\circ$ to $\delta = +2^\circ$, and has a cut at low Galactic latitudes that is designed to remove the highest stellar density regions.

2. The original footprint, but with no Galactic latitude cut.

3. A declination cut $-72^\circ < \delta < +12^\circ$, with the original cut at low Galactic latitudes.

4. A declination cut $-72^\circ < \delta < +12^\circ$, with a cut in Galactic extinction at $E(B-V) < 0.15$, or a reasonably smooth approximation to this. This could be crudely

approximated as a cut in Galactic latitude, e.g., $|b| < 15°$, as has been suggested in several white papers, but it is worth making a cut on more principled arguments.

Not included in the above is the option of the extended declination cut with no cut at low Galactic latitudes; we are assuming (but it should be checked) that there is no way to get to required LSST coadded depths uniformly over that full region. Note that we endorse experiments with different cadences at low Galactic latitudes; these are described in § 4.3.2 below.

4.2 The WFD Cadence

Cadence in WFD needs to be defined on a variety of timescales:

- The nature of individual visits (e.g., should they consist of two back-to-back 15-second snaps or a single 30-second exposure?)
- Returns to a given field in a given night: which filter(s) should be chosen?
- Returns to a given field within a given season.
- "Rolling" the cadence, in which in a given year, only a fraction of the available sky is observed, with an increased cadence in that year.

Before discussing each of these in turn, we note that for any given footprint and cadence, it may be that there is adequate time to significantly exceed the minimal coverage of WFD as defined in the LSST science requirements document (825 visits per field, summed over the six filters). For example, in the previous baseline OpSim run, there were an average of 910 visits per field over the WFD footprint, which took a total of 85% of the available observing time. Thus the simulations that follow can be done in two modes:

- Reaching the minimum number of total visits in WFD, and devoting remaining time to the mini-surveys, deep drilling fields, and other ancillary programs.
- Setting a firm limit (of perhaps 10%) for all ancillary programs, and maximizing the total visits or depth in WFD.

The SAC recommends that both these types of experiments be carried out in what follows.

It has already been demonstrated that the uniformity of the LSST sky coverage is much improved if the field centers are not fixed, but are dithered (in position and perhaps also rotation angle) by an appreciable fraction of the field size. We recommend that dithering be used for all OpSim experiments throughout the LSST footprint (both WFD and the mini-surveys described in § 4.3). Such large dithering is also needed for accurate flat-fielding and sky subtraction, as has been demonstrated with data from the Hyper Suprime-Cam survey (Aihara et al. 2018). Dithering is also an issue for the Deep Drilling Fields; see the discussion in § 4.4.

4.2.1 One or Two Exposures per Visit; the Length of a Visit

An LSST visit is defined to be 30 seconds on the sky. The arguments for that are compelling, and are outlined in the LSST overview paper and the science book, and we do not suggest varying this (except perhaps in the u-band; see below). The default has been to split this into two 15-second exposures ("snaps") on each visit, which may allow exploration of very short timescale variability and may make the rejection of cosmic rays and other artifacts more robust, as well as increase the dynamic range of LSST data at the bright end. However, carrying out a single 30-second snap provides two significant advantages, as pointed out in several white papers:

- By reducing the overhead of a second readout, the observing efficiency increases substantially, gaining back about 7% of the telescope time. This is of the same order as the total amount of time devoted to Deep Drilling Fields, and thus represents a substantial gain to overall LSST productivity.
- In 15 seconds, some of the exposures would be read-noise limited, especially in the u band under dark skies. 30-second exposures do not have this limitation, increasing the depth of the u-band observations by a few tenths of a magnitude.

While several white papers specifically requested the 2×15 second snaps, the science drivers for this were often relatively narrow, and would require processing beyond the current requirements on the LSST pipelines (which are planning simply to combine the two snaps). We therefore suggest that the default in the OpSim runs be a single 30-second exposure per visit, with a limited number of fiducial experiments run with the 2×15 second snaps[1].

The final decision on which of the two to run will depend on the as-delivered performance of the camera and instrument, as well as the image processing software. In particular, if it turns out that a pair of exposures is required to adequately remove cosmic rays and other artifacts, or the saturation and scattered light from bright stars give rise to much worse systematics with a single 30-second exposure, the project may have to use the two-snaps-per-visit model. Two snaps per visit also allows motions of fast-moving objects (e.g., near-earth asteroids and earth-orbiting satellites) to be determined from a single visit, although it is our understanding that this is not a limiting concern for determination of asteroid orbits.

Some of the white papers made suggestions for exposure times other than 30 seconds per visit. The most compelling of these were to do photometry over the entire sky that LSST would reach (including the extensions to the North, South, and at low Galactic latitudes) with significantly shorter exposures: 1 second and 5 seconds, in particular, to extend the LSST photometric system to brighter magnitudes. No specific case was made to do this repeatedly; there are many other smaller survey telescopes that are exploring the variable sky brighter than 16th or 17th magnitude (the nominal saturation limit of LSST). We

[1] Several white papers suggested that the two snaps of a visit be of different exposure lengths. There is no difference in the outputs of OpSim (other than directly calculable changes in depth due to read-noise) if we were to do this, so we do not suggest any specific experiments along these lines.

therefore recommend that the simulations include observations of the entire LSST footprint (including low Galactic latitudes) with 1 second exposures and 5 second exposures, in all 6 filters. The 1 second exposures can be done during twilight (§ 4.3.4). This can be done at any time during the ten years. We recommend that the entire LSST footprint be covered twice with one-second exposures in each filter, and twice with five-second exposures. once.

We also recommend experiments in which the u-band exposure time is varied. As noted above, with 15-second exposures, the u-band images will be read-noise-limited in dark skies, and the read-noise effects may not be negligible even in 30-second exposures, depending on the performance of the chips. It may therefore be advantageous to have u-band exposures longer than the canonical 30 seconds.

Yet another variant would be to adjust the exposure time of a given visit to give uniform depth: thus exposures in cirrusy conditions with a bright sky, high airmass and poor seeing would receive longer exposure times than the default 30 seconds. No white paper specifically recommended this strategy, and given that one has the freedom at the end of the survey to combine repeated observations in a variety of ways to reach a given depth, the SAC did not see a strong argument to explore this option. OpSim does include metrics that quantify the uniformity of depth across the survey area, and results to date indicate that over the WFD footprint, the coadded depth is quite uniform. However, if WFD is extended further North and South, as some of the experiments above recommend, this uniformity may not hold up, and it may be worth exploring this option.

4.2.2 The Choice of Fields and Filters on a Given Night, and in a Given Month

The LSST cadence requires repeat visits to a given field in a given night to measure "tracklets", i.e., the proper motion/parallax vectors of solar system asteroids, which are linked together to determine orbits. We assume (as did all the white papers we received) that a pair of visits is adequate to measure a tracklet, and to distinguish asteroids from (stationary) transient events and artifacts; this assumption is based on the expected performance of the Moving Objects Processing System (MOPS, which will determine asteroid orbits) and the flagging and rejection of instrumental artifacts.

There is freedom in deciding in which band these repeat visits happen. Some white papers requested that exposures in r be repeated in r (as this maximizes sensitivity to faint red Kuiper-Belt Objects) while others emphasized the need for pairs of g and i exposures on a given night, to optimize the detection of tidal disruption events around supermassive black holes in galaxy nuclei. A particularly interesting suggestion is made by Bianco et al, "*Presto-Color*: An LSST Cadence for Explosive Physics & Fast Transients", which includes three observations per night in two different filters, as a way to identify and characterize fast transients. No one of these observing modes need be carried out through the whole ten-year survey, and there is a strong argument for using different modes at different times. With all this in mind, we recommend the following sets of rules governing the repeat exposures to explore these different options:

- There is no restriction on the filters in which the second visit happens.

- The second visit is forced to be in a different filter from the first visit (to allow colors to be measured for all variable sources).
- If the first visit of the night of a given field is in r, the second visit should be in r as well. Each field should be revisited at least six times in a given season. (This may happen by default with any cadence; it is certainly likely to happen with rolling cadence). If the first visit of the night of a given field is in g, the second visit should be in i. A repeat observation of the same field in g the following night is then given high priority.
- The Presto-Color option: pairs of visits are carried out within 0.5 hour in a pair of filters (ideally g & i or r & z), followed by another observation in one of those filters later in the night. The requirement on three visits in a given night, with two of them within 0.5 hours, is restrictive enough that this should be carried out over only a modest fraction of the WFD survey. It may be that this cadence can more naturally be accommodated in DDF observations.

Comparing these experiments will determine how restrictive the filter choice on repeat visits is in practice, and whether it has any sort of deleterious effect on LSST observing efficiency or survey coverage uniformity.

There were no strong requests for repeat observations in u or y. Given that the LSST camera can hold only five of the six filters at any given time, we recommend that the u-band filter observations be concentrated within ± 2 days of New Moon. This will allow the y-band to be swapped into the filter wheel for the rest of the month; observations in y can take place when the moon is up, as well as during twilight.

At least at high Galactic latitudes, most of the variability science described in the white papers was focused on the bluer filters; there were no calls for specific cadence in the z and y filters.

One white paper (K. Bell et al., "A Cadence to Reduce Aliasing in LSST") makes a specific suggestion about offsetting the time of night a given field is observed in sequential nights to minimize the 24-hour aliasing. This is a good idea, but will require developing new code within OpSim to implement, and experiments with OpSim to understand its impact on observing efficiency. A metric is also needed to demonstrate whether the anti-aliasing is effective. This could be as simple as making a fiducial star with a sinusoidal light-curve with a 1-day period, and checking how well it gets fit at the end of the survey. These experiments should be put at high priority; if it is shown that this can be made workable, and that the anti-aliasing is as effective as it claims to be, this approach should be made the default for the OpSim runs to follow.

4.2.3 Rolling Cadence and the Length of a Season

LSST is often described as able to cover the full observable sky at any given time in three nights (in a single filter). That is the default in the baseline OpSim runs; the area of sky covered at any given time is maximized. Given the multiple LSST filters, this means that the time between visits in a given field in a given filter is often of the order of two weeks in the current LSST baseline cadence. A number of white papers argued strongly that a more rapid cadence is needed; this can be done by using a so-called *rolling cadence*: in any given year, the area of sky to be observed should be a fraction of the full area available, allowing a higher cadence for that year (in exchange for concentrating in a different area of sky in following years). There are clearly trade-offs: there is a compelling need to observe the full LSST footprint at at least moderate depth each year in all six filters (e.g., for measurements of very long timescale variability, as well as proper motions), and putting too much emphasis on rolling will affect the uniformity of LSST coverage at intermediate data releases.

We suggest the following cadence approaches for all four Wide-Fast-Deep footprints described in § 4.1.

- The universal cadence for all 10 years (which has defined the baseline cadence).
- Split the WFD sky in two equal-area halves, separated by declination. Observe each half in alternate years, but obtain modest coverage over the remaining footprint.
- As the above, but devoting the first and last years of the survey to full-footprint observations. Given the way sky is covered through the seasons, getting uniform coverage over the LSST footprint will require a half-year of transition from full-footprint to rolling mode, after the first year; there would need to be a similar half-year of transition starting in Year 8.5, before returning to full-footprint mode in the last year of the survey.
- Split the sky into three equal parts, defined by declination. There are two ways this could be observed. The first would be to roll in the six years between years 1.5 and 7.5, with the period between 0 and 1.5 years, and between 7.5 and 10 years, devoted to full-footprint observations. The second would be to devote Years 1, 5, 9, and 10 to full-footprint observations. In Years 2, 3, and 4, and Years 6, 7, and 8, focus observations in one-third of the sky each year. In each of these scenarios, get a modest number of visits in each filter over the full footprint each year.
- Similar to the previous option: split the sky into *six* equal parts, and roll in the six years in between Years 1.5-7.5, with the remainder of the time devoted to full-footprint observations.
- The science case for splitting up the sky in more parts is less clear. The OpSim team is able to simulate up to a 10-way split, but the SAC did not find a compelling reason to, in effect, do a different area of sky each year of the survey.

In the rolling experiments above, we have suggested that each year, we cover the full WFD footprint with a modest number of visits. This can be specified in OpSim in terms of a scale-down factor relative to the universal cadence number of visits per year. We suggest experiments in which this scale-down factor is 0.2, 0.1 and 0.05.

The rolling suggested here requires a way to divide up the WFD footprint. We have suggested doing this with declination cuts, but as the visibility of a given area of sky through the night and through the year is a function of declination, it is possible that is not the ideal way to do so. None of the white papers had specific suggestions about how the sky should be divided up.

A number of white papers expressed a desire to maximize the length of a season on any given area of sky. That is, in a given area of sky in a given year, one would like to observe for as many months as that area is at reasonable airmass, to measure supernova light curves and other variable phenomena of several-month duration. The current implementation of OpSim approximately maximizes the season length given restrictions on airmass, but it would be useful to run an experiment to test this, namely a rolling in which extra weight is given to visits that extend an observing season, to see what effect this has on metrics that capture the sensitivity of the survey to long timescale variable phenomena. Note that for a given number of visits in a given year, decreasing the length of a season reduces the inter-night spacing, allowing short-timescale variability to be better characterized.

4.3 Mini-Surveys: Extensions to the North, South, Low Galactic Latitudes, and Twilight

In the context of this document, a mini-survey is defined as an area of sky larger than a single pointing, that is observed with a cadence different from that of WFD. There is not enough observing time to incorporate all of these mini-surveys in OpSim experiments, especially with the more expansive requested footprints, and it is work for the OpSim team to determine which combinations can be done. Fewer observations in mini-surveys mean more time for the WFD survey and the DDFs. To understand the tradeoff, we recommend a series of experiments in which, e.g., the South Celestial Pole is not covered, and the time freed up is devoted to WFD.

4.3.1 Pushing North

The default footprint for WFD extends to $\delta = +2°$; in § 4.1, we suggest experiments in which this limit is extended to $\delta = +12°$. Several white papers made the case for mini-surveys that extend the LSST footprint further North yet, to $\delta = +30°$, in order to maximize the overlap with other major surveys (especially DESI and Euclid), and to maximize the coverage of the ecliptic plane. With these arguments in mind, we suggest the following mini-surveys:

- In combination with the standard WFD footprint, extending to $\delta = +2°$, observe in *ugrizy* a total of $3, 10, 10, 3, 3, 3$ visits, respectively over the declination range $+2° < \delta < +30°$. To the extent possible, distribute these visits in time following the recommendations of Table 6 of the Capak et al. white paper, "Mini-Survey of the Northern Sky to Dec $< +30°$."

- In combination with the extended WFD footprint, extending to $\delta = +12°$, observe to the same depth as described above in the declination range $+12° < \delta < +30°$.
- In both of the above cases, increase the number of visits to 15, 10, 10, 5, 21, 3, respectively, if the time is available.

As a separate experiment, carry out the Northern Ecliptic Spur experiment as described in the white paper by Schwamb et al., "A Northern Ecliptic Survey for Solar System Science". This involves observing the region within $\pm 10°$ of the ecliptic North of the WFD footprint. Each field is to be observed on six nights per month for a five-month period (with pairs of exposures per night), over 7 of the LSST years. In the remaining 3 years, there would be a total of 15 visits per field spread over 2-3 months. These observations would be done in g, r and z, with 60% of the observations in r.

4.3.2 Pushing South

There is a compelling case to extend the LSST footprint to the South Celestial Pole, in part because of the presence of the Magellanic Clouds. The experiment described in the white paper by Olsen et al., "Mapping the Periphery and the Variability of the Magellanic Clouds," should be included. In particular:

- A survey of the entire region South of the Southern limit of the WFD survey: 40 30-second visits per field in each of *ugrizy*.[2] This should be done both for the standard footprint, extending to $\delta \sim -62°$, and for the extended footprint, which goes to $\delta = -72°$.
- Separately, a total of 2000 15-second exposures in each of 3 pointings in the center of the SMC, and 9 pointings in the center of the LMC, as described in the above white paper. These would consist of $50, 1300, 300, 300, 30$ and 20 visits in *ugrizy*, respectively. Note that some of these fields are in the extended WFD footprint; the above number of exposures would include visits that are part of WFD. We also suggest an experiment in which the number of visits in g is reduced from 1300 per field to 300 per field. Note that these visits are half as long as for the main survey; the cadence is described in the above white paper.

4.3.3 Low Galactic latitudes

We received a variety of white papers with various suggestions for covering the full low-galactic-latitude footprint. As we mentioned in § 4.1, there are multiple possible definitions of what constitutes low Galactic latitudes. The science case for low latitudes is mostly focused on the bulge region. Different white papers define the extent of the bulge differently;

[2] Olsen et al. specifically excluded the y band in their request, but we recommend that it be included in the OpSim experiments; there will likely to be other science investigations using these data that will take advantage of the y-band photometry.

here we suggest a strawman definition of $-20° < l < +20°, -10° < b < +10°$, where l and b are Galactic coordinates. It is only in the bulge that one reaches the crowding limit well before LSST's full WFD depth (but variability science will of course benefit from repeated observations in the bulge). The scientific case for observations of the MW anti-center has not been stated strongly; the only specific request (from the white paper by Prisinzano and Magrini) was for WFD-depth observations at longer wavelengths, to identify pre-main sequence M stars in young open clusters.

With this in mind, we recommend experiments in which the low-latitude region (as defined as the relevant complement of the WFD footprint, as described in § 4.1) is split into two subregions:

- The bulge region, as defined above.
- The rest of the low latitude sky. This would not extend northward of the northmost limit of the WFD footprint ($\delta = +2°$ and $\delta = +12°$ in the scenarios described in § 4.1); that is, there would be no spur of low-latitude sky sticking above the main survey, as was done in the baseline OpSim runs carried out in the past.

Here are suggestions for the cadence for those fields:

- Following the Olsen et al. "A Big Sky Approach to Cadence Diplomacy" white paper, both the bulge and the rest of the low-latitude region receive 250 visits over ten years. As a default, the number of visits in each filter would be in the same ratio as in the WFD footprint. This would be "rolled" in the same way that WFD would be rolled, as described in § 4.2.3.
- Adopt the WFD depth and cadence in the bulge, and up to 250 visits for the rest of the low-latitude sky (limited perhaps by the total telescope time available).
- As above, devoting 825 visits to each field in the bulge, but doubling the number of i-band visits (and reducing the number of visits in the other bands) to study proper motions in the bulge (see the Gonzalez et al. white paper, "The Definitive Map of the Galactic Bulge").
- As above, but requiring a visit every 3 days in one of *griz* to every bulge field, when it is available. This is inspired by the white paper by Street et al. "The Diverse Science Return from a Wide-Area Survey of the Galactic Plane," to look for transits and microlensing in the bulge.

Note that we have already suggested a definition of the WFD footprint with no low-latitude cut at all. Thus by definition, all low-latitude regions would receive a WFD cadence. A number of white papers also suggested observations with shorter exposures (5 seconds) to increase dynamic range in crowded regions; a pair of visits with short exposures is already called out in § 4.2.1 for the full LSST footprint.

4.3.4 Twilight observing

On any given night, LSST observing happens between 12 degree twilight in the beginning and end of the night. While it may be possible to make additional observations in brighter twilight time (i.e., sun higher than -12 degree altitude), the sky background will be much higher and difficult to model in OpSim. We received two white papers with recommendations for mini-surveys to be carried out during twilight. We recommend a pair of experiments: one in which these mini-surveys are in competition for the WFD observations at any time of the night, and one in which they are restricted to happen between 18 and 12 degree twilight. Despite the difficulties of modeling the sky brightness between 6 and 12 degree twilight, it may also be useful to include simulations that use that time for the survey for near-earth asteroids below.

The two twilight mini-surveys are:

- Seaman et al., "A near-Sun Solar System Twilight Survey with LSST." This proposal seeks to maximize repeat observations in twilight to find near-earth asteroids and constrain their orbits. There is a specific cadence of observations described in the white paper. Given the color of the twilight sky, these observations should be done at longer wavelengths, i.e., $rizy$.
- Richards et al, "Leveraging Differential Chromatic Refraction in LSST." This aims for a single visit in u and g over the entire sky at large airmass (> 1.5). The only reason this was suggested to be done in twilight is because this time is in less direct competition with the main survey, but the S/N in these bluer bands will be affected by the bright twilight sky. However, the total time requested is modest, and thus experiments in which this program is carried out during dark time should also be done.

4.4 Deep Drilling Fields: Location and Cadence

The LSST project has long agreed that four so-called Deep Drilling Fields (DDF), each the size of a single LSST pointing, would be observed more often than the main survey; the default is to cover them with enough visits to go roughly one magnitude deeper in each of the filters than in the uniform WFD survey. The cadence optimization white papers included a number of ideas for additional DDFs, as well as suggestions on what the cadence in the DDFs should be.

The DDFs will need to be dithered (perhaps both in boresight position and rotation angle), in order to:

- make the coverage over the DDF footprint more uniform, especially with regard to chip gaps;
- allow correction for artifacts such as bad columns, hot pixels, and other CCD effects;

- allow for accurate sky subtraction and flat-fielding.

The first two of these concerns can probably be ameliorated with relatively small dither offsets, of the order of a single chip, while the third of these requires dithers an appreciable fraction of the full LSST pointing. The DDFs all fall within the WFD footprint anyway, and so will by default already include visits in the WFD mode with large dithers (see § 4.2); depending on the nature of scattered light in the as-built LSST camera, and the data processing software, the fraction of the observations in the DDF fields with large dithers may be adequate to do proper sky subtraction. Until we can verify that this is the case, it will be useful to carry out OpSim runs with both a modest dither (of order a chip or so) and a large dither (of order half the field size).

Only a few white papers made specific suggestions for cadence in the default DDFs. The two that should be experimented with are the following:

- The cadence described in Scolnic et al., "DESC Recommendations for the Deep Drilling Fields and other Special Programs." This is designed to allow multi-band light curves of supernovae to be measured. Reminiscent of the rolling cadence suggestions (§ 4.2.3), this calls for observing a given field for long seasons, more than six months, interweaving gri and zy observations every 3 days with 2, 4, 8, 25, 4 visits in $grizy$, respectively. This white paper makes no specific recommendation for u-band cadence.
- The cadence described in Brandt et al., "Active Galaxy Science in the LSST Deep-Drilling Fields." This calls for repeat observations every two days in each of $grizy$, again maximizing the length of the observing season on each field; the somewhat denser cadence here is important, e.g., for studies of AGN accretion disks using continuum reverberation mapping. In this case, the observations every two days would have 1, 1, 3, 5, 4 visits in $grizy$ (making this program considerably less expensive in total than the DESC DDF program). This white paper also stresses the importance of u-band coverage for active galaxy science, with the best time-domain coverage that LSST can provide given practical filter-wheel constraints (the equivalent of 4 u-band visits every two days is desired).

The suggestions in these two white papers could also potentially be mixed together into a hybridized DDF program. For example, the preferred AGN cadence could be adopted in some years, and the preferred DESC cadence could be adopted in others. In such possible hybridized programs, the depth requirements of one program will likely need to be respected when the other program is setting the cadence (e.g., deep $griz$ for supernova studies, and u-band coverage for AGN studies). Also, the choice of the specific years when one approach is dominant should consider related ongoing observations (e.g., spectroscopic AGN reverberation mapping with SDSS-V Black Hole Mapper and 4MOST TiDES-RM). Simulations of possible hybridized programs are encouraged, in consultation with the relevant white-paper authors.

Both the DESC and AGN DDF white papers provide specific suggestions for the precise locations of each DDF, and these are in agreement. These precise locations should be adopted for the simulations. Regarding the XMM-LSS field, consideration of possible scattered-light effects due to the nearby bright star Mira remains important, as detailed in the Brandt et al. white paper. There is also agreement in the white papers that additional DDF observations during LSST commissioning are important.

Several white papers also argue for observations in "movie mode", in which a given field is observed continuously in a single filter as long as it is up in a given night, to study variability. This would be an experiment to carry out in a single field on a single night, and could be done during LSST commissioning; this does not require a separate OpSim experiment.

Given that the LSST has already committed to four DDFs, there is limited opportunity for additional DDFs. We have already made reference to special fields in the Galactic Bulge and the Magellanic Clouds in §§ 4.3.3 and 4.3.2, respectively. We suggest exploring one other DDF[3], centered roughly at $\alpha, \delta =$ 04:44:00, −53:20:00, as suggested both by Scolnic et al: , and Capak et al in "Enhancing LSST Science with *Euclid* Synergy". The Euclid request asks for a total of 20 deg^2 (i.e., two LSST pointings), whose exact position is still being optimized by the Euclid team. They require only modest depth, comparable to the WFD co-added depth in *gri*, and a magnitude fainter in *uzy*. However, they do request that these fields be prioritized for the first few years of LSST observations, to maximize the scientific overlap with LSST.

There were other suggestions for DDFs selected in coordination with WFIRST, both at high latitudes (to observe supernovae) and low (for microlensing). These would require coordination with WFIRST both in position on the sky, and in time, to interleave the observations with the two facilities. Ideally, the WFIRST field-of-regard should be improved to allow better overlap with some of the already planned LSST DDFs. In particular, even a modest improvement in the WFIRST field-of-regard should allow the Chandra Deep Field-South and the ELAIS-S1 field to be used as WFIRST supernova fields, greatly improving efficiency from an LSST perspective and also improving the available multiwavelength data. Further work is needed to develop the relevant capabilities within OpSim, and to specify exactly what the LSST cadence would be.

Trilling et al. ("Deep Drilling Fields for Solar System Science") suggest a series of observations on five fields along the ecliptic, spread out in ecliptic longitude (and at high Galactic latitude, to minimize the stellar density), to measure the orbits and colors of very faint Kuiper Belt Objects. Each field would be observed on four different nights, for 2.1 hours on each night. On a given night, the field would be observed in a single filter (in movie mode with standard 30-second exposures). The second night on a given field would be 1-2 months after the first night, the third night would be the following year, and the fourth visit would be the year after that. Three of the four nights would be in r with the fourth night in g.

If time allows, it would be worthwhile to try experiments in which the four standard

[3] Several of the DDFs suggested in the white papers either were pointed at objects much smaller than LSST's field of view, thus more suited for observations with other facilities, or involved cadences that would lead to highly inefficient observing.

DDFs, the Deep Field recommended by DESC and the Euclid team, and the five ecliptic plane fields are all observed. It is not yet clear whether the available observing time is adequate to do that in the expanded WFD footprint experiments described in § 4.1. Note further that the DESC and Euclid team requests are not identical, although they are in similar areas of sky; the Euclid team wants two adjacent pointings, to be prioritized earlier in the LSST survey, going deeper than the WFD survey only in uzy, while the DESC request is a single pointing, going deeper in all six bands.

4.5 Specialized Observing Modes

We received a number of white papers advocating for specific target-of-opportunity observations. The SAC strongly endorses the idea that LSST have a mechanism to respond to scientifically compelling targets of opportunity. In practice, a target of opportunity observation would interrupt the normal observing mode of LSST, and lead to a modest number of exposures that may or may not be incorporated into the main survey. The SAC suggests that the interruptions due to a target-of-opportunity observation be considered like bad weather, taking a small amount of time away from the main LSST survey.

The white paper by Margutti et al, "Target of Opportunity Observations of Gravitational Wave Events with LSST", estimates that effective follow-up of LIGO-Virgo events with LSST will require roughly 85 hours per year, about 2% of the total available time. The SAC recommends that an OpSim experiment incorporate ToOs from gravitational wave triggers as outlined in Margutti et al. We note that this white paper was written prior to LIGO-Virgo's third observing run. Given the success of the first month of this new run, including two neutron star merger alerts within 30 hours, estimates of the kilonova rates, luminosity, timescales, colors, and positional uncertainties could change substantially over the next year. The SAC recommends re-visiting the assumptions made in this white paper in six to twelve months.

One white paper, Thomas et al, "Unveiling the Rich and Diverse Universe of Subsecond Astrophysics through LSST Star Trails," suggests a specific experiment in which images are trailed to look for very short timescale variability. This would take a small amount of time, and the resulting data will not contribute directly to the LSST main survey. We therefore do not suggest an OpSim experiment for this.

As we have discussed elsewhere in this document, there were several white papers that asked for observations synchronized with other facilities, in particular WFIRST and Euclid. We look forward to working with the OpSim team to figure out the best way to simulate such observations and to assess their scientific potential and impact on the rest of the survey.

5 Concluding Thoughts

As we have described in this report, a given OpSim run is specified by:

- The assumed footprint of the Wide-Fast-Deep survey; we suggest four footprints in § 4.1, defined by their North and South limits and the definition of what constitutes low latitude.
- The cadence of the WFD observations, which involves decisions about the nature of a single visit, and how the sky is covered on a nightly, monthly and yearly basis.
- The choice of mini-surveys, including those that push further North, further South, and to lower Galactic latitudes than the WFD survey;
- Deep Drilling fields and their cadence.
- Additional observing programs, such as Target of Opportunity and observations coordinated with other surveys.

In the above, we have outlined our recommendations for experiments in each of these categories. As we have emphasized, not all combinations of the various tests result in a survey that fits within the ten-year survey, and satisfies the core LSST science requirements. The SAC is working with the OpSim team to determine the combinations that make sense.

Understanding how all the different choices for an OpSim run affect each other and the science outputs of the survey is a complex problem. However, not every combination needs to be tried; some of the experiments we suggest can be done on a modest number of fiducial OpSim runs. For example, we know that there is a roughly 7% hit on observing efficiency if the standard visit consists of two 15-second snaps rather than a single 30-second exposure. We recommend carrying out the bulk of the OpSim runs with the latter, with a modest number of runs using the two 15-second snaps, to quantify the effect on scientific outputs. In the same spirit, we suggest a range of different choices for selecting the filter for the second exposure of a given field on a given night (§ 4.2.2); the effects of these choices on overall observing efficiency and science output can be determined with a few focused experiments. As we have emphasized earlier, this will be an iterative process: as the OpSim team carries out experiments and makes the outputs available to the LSST scientific community, we will collectively learn what works well and what does not, informing the next round of experiments.

A key tool in all this work will be the development of metrics to be run on OpSim outputs, to assess the scientific impact of each of these choices. Without quantitative metrics, it will not be possible to understand the tensions between different science programs as the survey strategy algorithms are modified. The OpSim team may need additional resources to write those metrics in the Metric Analysis Framework.

The SAC and the OpSim team are eager for feedback on all topics discussed above, especially as the OpSim outputs become available. As indicated earlier, we recommend that `https://community.lsst.org/c/sci/survey-strategy` be used for discussion, and specific recommendations via a google form at `https://tinyurl.com/OpSim-feedback`.

Appendix: The submitted white papers

The 46 white papers are listed below, grouped thematically in rough correspondence to the order in which different topics are discussed in § 4. The first or corresponding author is listed in each case. The white papers are available publicly at: `https://www.lsst.org/submitted-whitepaper-2018`.

Wide-Fast-Deep Footprint

- **Knut Olsen**: A Big Sky Approach to Cadence Diplomacy
- **Michelle Lochner**: Optimizing the LSST Observing Strategy for Dark Energy Science: DESC Recommendations for the Wide-Fast-Deep Survey

Pushing North or South; Synergy with Other Surveys

- **Peter Capak**: Mini-survey of the northern sky to Dec $< +30$
- **Adam Bolton**: Maximizing the Joint Science Return of LSST and DESI
- **Peter Capak**: Enhancing LSST Science with Euclid Synergy
- **Colin Snodgrass**: Simultaneous LSST and Euclid observations — advantages for Solar System Objects
- **Meg Schwamb**: A Northern Ecliptic Survey for Solar System Science
- **Chris Hirata**: Enabling Cosmological Synergies between the LSST and WFIRST High Latitude Survey
- **Ryan J. Foley**: LSST Observations of WFIRST Deep Fields
- **Rachel Street**: Unique Science from a Coordinated LSST-WFIRST Survey of the Galactic Bulge

Low Galactic Latitudes

- **Jay Strader**: The Plane's The Thing: The Case for Wide-Fast-Deep Coverage of the Galactic Plane and Bulge
- **Mike Lund**: A Higher Cadence Subsurvey Located in the Galactic Plane
- **Oscar Gonzalez**: The Definitive Map of the Galactic bulge
- **Loredana Prisinzano**: Investigating the population of Galactic star formation regions and star clusters within a Wide-Fast-Deep Coverage of the Galactic Plane
- **Rachel Street**: The Diverse Science Return from a Wide-Area Survey of the Galactic Plane
- **Massimo Dall'Ora**: unVEil the darknesS of The gAlactic buLgE (VESTALE)

Cadence Decisions in WFD

• **Keaton Bell**: A Cadence to Reduce Aliasing in LSST

• **Gordon Richards**: Testing of LSST AGN Selection Using Rolling Cadences

• **Federica Bianco**: Presto-Color: An LSST Cadence for Explosive Physics & Fast Transients

• **Suvi Gezari**: A Smart and Colorful Cadence for the Wide-Fast Deep Survey

• **Katja Bricman**: TDEs with LSST

• **Igor Andreoni**: A strategy for LSST to unveil a population of kilonovae without gravitational-wave triggers

• **Kat Volk**: The Effects of Filter Choice on Outer Solar System Science with LSST

• **Mario Juric**: Enabling Deep All-Sky Searches of Outer Solar System Objects

• **Aprajita Verma**: Strong Lensing considerations for the LSST observing strategy

• **Claudia Raiteri**: Blazars and Fast Radio Bursts with LSST

• **Roberto Silvotti**: Searching for white dwarf transits with LSST

Deep Drilling Fields and Mini-Surveys

• **Dan Scolnic**: Optimizing the LSST Observing Strategy for Dark Energy Science: DESC Recommendations for the Deep Drilling Fields and other Special Programs

• **Niel Brandt**: Active Galaxy Science in the LSST Deep-Drilling Fields: Footprints Cadence Requirements, and Total-Depth Requirements

• **Benne W. Holwerda**: Large Synoptic Survey Telescope White Paper; The Case for Matching U-band on Deep Drilling Fields

• **Keaton Bell**: Continuous Cadence Acquisition of the LSST Deep Drilling Fields

• **Eric Feigelson**: Characterizing Variable Stars in a Single Night with LSST

• **Gisella Clementini**: The Gaia-LSST Synergy: resolved stellar populations in selected Local Group stellar systems

• **Rosaria Bonito**: Young Stars and their Variability with LSST

• **Eric Bell**: LSST Local Volume mini-survey

• **David Trilling**: Deep Drilling Fields for Solar System Science

• **Knut Olsen**: Mapping the Periphery and Variability of the Magellanic Clouds

• **Radek Poleski**: The First Extragalactic Exoplanets — What We Gain From High Cadence Observations of the Small Magellanic Cloud?

Specialized Observing Modes and Targets of Opportunity

• **John Gizis**: Calibrating Milky Way Maps: An LSST Bright(ish) Star Survey

• **Gordon Richards**: Leveraging Differential Chromatic Refraction in LSST

• **Rob Seaman**: A near-Sun Solar System Twilight Survey with LSST

• **David Thomas**: Unveiling the Rich and Diverse Universe of Subsecond Astrophysics through LSST Star Trails

• **Chris Walter**: LSST Target of Opportunity proposal for locating a core collapse supernova in our galaxy triggered by a neutrino supernova alert

• **Graham Smith**: Discovery of Strongly-lensed Gravitational Waves – Implications for the LSST Observing Strategy

• **Raffaella Margutti**: Target of Opportunity Observations of Gravitational Wave Events with LSST

• **Seppo Laine**: LSST Cadence Optimization White Paper in Support of Observations of Unresolved Tidal Stellar Streams in Galaxies beyond the Local Group

Membership of the Science Advisory Committee

Timo Anguita (Universidad Andres Bello, Chile)
Franz Bauer (Universidad Católica de Chile)
Rachel Bean (Cornell University)
Niel Brandt (Penn State University)
Mansi Kasliwal (Caltech)
David Kirkby (University of California Irvine)
Charles Liu (CUNY College of Staten Island/AMNH)
Amy Mainzer (Jet Propulsion Laboratory)
Meg Schwamb (Gemini Observatory)
Joshua Simon (Carnegie Observatories)
Stephen Smartt (Queen's University Belfast)
Anže Slosar (Brookhaven National Lab)
Michael Strauss (Princeton University; chair)
Lucianne Walkowicz (Adler Planetarium)
Risa Wechsler (Stanford University)

The Science Advisory Committee acknowledges the help, assistance and advice of many people, including: Lynne Jones, Peter Yoachim, Tiago Ribeiro, Željko Ivezić, and the LSST Data Management team. Most importantly, we thank the authors of the white papers, who contributed so many useful ideas to this whole process.

References

Aihara, H. et al. 2018, PASJ, 70, SP1, S4 (HSC)

Delgado, F. et al. 2014, Proceedings of the SPIE, 9150, 915015 (OpSim)

Jones, R.L. et al. 2014, Proceedings of the SPIE, 9149, 91490B (MAF)

Ivezić, Ž. et al. 2018, `https://docushare.lsst.org/docushare/dsweb/Get/Document-28382` (call for cadence optimization white papers)

Ivezić, Ž. et al., `https://docushare.lsstcorp.org/docushare/dsweb/Get/LPM-17` (SRD)

Ivezić, Ž. et al. 2019, ApJS, 873, 111 (LSST overview paper)

LSST Science Collaborations 2009, arXiv:912.0201 (LSST Science Book)

LSST Science Collaborations 2017, arXiv:1708.04058 (COSEP)