Draft version July 20, 2020
Typeset using LaTeX **modern** style in AASTeX62

Survey Strategy and Cadence Choices For the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST)

R. Lynne Jones,[1] Peter Yoachim,[1] Ribeiro Tiago,[2] and Željko Ivezić[1]

[1] *University of Washington, Dept. of Astronomy, Box 351580, Seattle, WA 98195, USA*
[2] *LSST Project Office, 950 N. Cherry Ave., Tucson, AZ 85719, USA*

(Dated: July 20, 2020)

ABSTRACT

A summary of survey strategy and cadence choices, simulated and evaluated by the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST) Scheduler Team, prepared for the Survey Cadence and Optimization Committee (SCOC).

The initial idea of a large telescope survey, covering the entire visible sky repeatedly every few days in multiple bandpasses over the course of ten years, is the core idea of the LSST. A large area (about 20,000 square degrees) observed under a wide range of conditions to deep coadded limiting magnitudes in bandpasses *ugrizy* enables cosmological studies with unprecedented precision; the same survey, when cadenced well, can serve to open new windows into our understanding of transient and variable stars, and extend our knowledge of small bodies throughout the Solar System by orders of magnitude. The outlines of these goals and some basic necessary requirements for those goals are outlined in the LSST Science Requirements Document (SRD)[a]. Finding options for the survey strategy to meet more detailed needs of an even wider range of science goals, as well as building the LSST Scheduler and Metrics Analysis Framework, has been the work of the LSST Scheduler Team with support and input from the astronomical community, including the COSEP [b], the Call for White Papers[c], and innumerable metrics, and guidance from the LSST Science Advisory Committee in their Recommendations for Operations Simulator Experiments[d].

1. INTRODUCTION

Note: This paper needs to focus on survey strategies and their evaluation.

Introduction - cover basic idea of survey simulator, scheduler and weather/telescope models.

[a] ls.st/srd
[b] https://github.com/LSSTScienceCollaborations/ObservingStrategy
[c] Document-28382
[d] Document-32816

Cover basic survey strategy starting point - wide area, frequent coverage, ten year timespan - and why. xxx–scheduling LSST is a balance between meeting the SRD requirements, and optimizing the return of the four primary science goals.

Mention COSEP and call for white papers - idea is to do the best science we can, add last 10% "best" science.

Earlier attempts at simulating LSST in Rothchild et al. (2019) and Naghib et al. (2019).

2. SURVEY SIMULATOR OVERVIEW

Probably need some reference to what survey scheduler was used / how it was set up for various runs, how the runs were performed, and what the input weather and telescope models were like.

2.1. *The Model Observatory*

Discuss kinematic model, seeing model, weather model.

2.2. *The Scheduler*

The scheduler is designed to provide real-time decisions on where and how to observe. Because we expect there to be things like weather interruptions, we need a system that can recover quickly. Unlike other traditional telescope schedulers, we do not try to optimize a large number of observations in advance, but rather use a decision tree along with a modified Markov Decision Process. The scheduler behavior is set by a small number of free parameters that can be tuned.

Our baseline scheduler uses a three tier decision tree when deciding what observations to attempt.

2.2.1. *Tier 1: Deep Drilling Fields*

The first tier of the decision tree is to check if there are any deep drilling fields that should be executed. We typically have five DDFs in a simulation.

For a DDF to be eligible to send a sequence to the observing queue, it must

- Not currently be twilight
- Have enough time to finish a sequence before twilight begins
- Be in it's target hour angle range
- The moon must be down
- The DDF must not have exceeded it's limit of observations (typically ∼1% of the total number of visits)

If the DDF has not fallen behind, it will space sequences by at least 1.5 days. There is also a check to see if the DDF will be feasible and better observed later in the night, in which case no observations are requested.

Name	RA (Deg)	Dec (Deg)
ELAISS1	9.450	-44.000
XMM-LSS	35.708	-4.750
ECDFS	53.125	-28.100
COSMOS	150.100	2.182
EDFS	58.970	-49.280
EDFS	63.600	-47.600

Table 1. The location of the deep drilling fields used in our simulations.

If the above conditions are met, the DDF sends it's sequence of observations to the queue to be executed. There are currently no attempts at recovery if a sequence is interrupted.

The spatial position of the DDF is dithered nightly up to 0.7 degrees. The camera rotator is also varied nightly to be between -75 and 75 degrees with respect to the telescope.

2.2.2. *Tier 2: The Blobs*

If there are no DDFs requesting observations, the decision tree moves to the second tier. This tier is the survey workhorse, executing ∼80% of the simulation visits. This tier will only request observations if it is not currently twilight, and there is at least 30 minutes before twilight begins.

A modified Markov Decision Process (MDP) is used to decide what sky region and filter combination to observe given the current conditions and observation history. Briefly, the MDP balances the desire to observe areas 1) that are closest to the optimal possible in terms of 5-sigma depth, 2) which have fallen behind the specified desired survey footprint, 3) are near the current telescope pointing and 4) in the currently loaded filter to minimize filter changes. In addition to these core components, the MDP includes a mask around zenith, a 30 degree mask around the moon, and small masks around the bright planets (Venus, Mars, Jupiter). The end product of the MDP is a reward function that ranks the desirability of every point in the sky. Because this tier does not execute in twilight, we assume the reward function is relatively stable on 40 minute timescales.

A sky area around the reward function maximum that will take ∼22 minutes to observe (∼35 pointings) is then selected. If possible, the area is selected to be be contiguous. The exact position of the telescope pointings are determined by the sky tessellation, which is randomly oriented for each night. The camera rotator angle (relative to the telescope) is also randomized between ±80 degrees each night.

A traveling salesman algorithm is used to put the pointings in an order that minimizes the slew time. The list of pointings are then repeated, usually in a different filter, ensuring moving objects can be detected. One of seven possible filter combinations is used: $u + g$, $u + r$, $g + r$, $r + i$, $i + z$, $z + y$, or $y + y$. We use 30 second visits

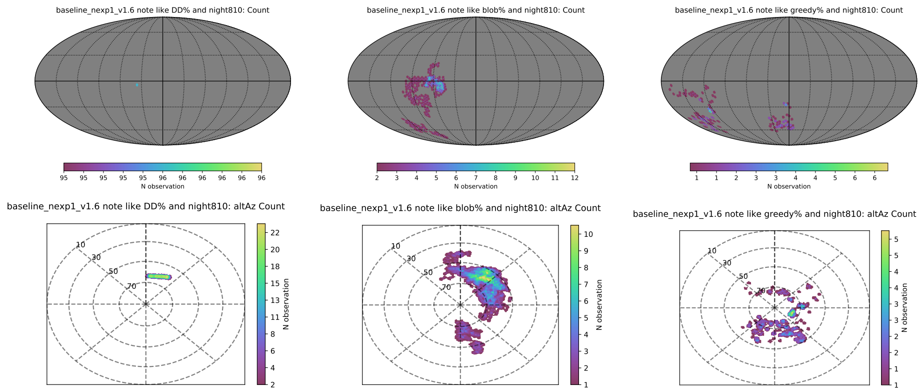

Figure 1. Examples of how the three scheduler tiers execute during a single night. Left panels show how a DDF sequence was observed during the night. Middle panels show observations taken as part of blob pairs. Right panels show the greedy observations taken in twilight time.

for the majority of simulations. The official baseline uses visits comprised of two 15 second snaps.

2.2.3. *Tier 3: Greedy*

If it is during morning or evening twilight, or close to morning twilight, the DDFs and Blob surveys will pass and the decision tree goes to the third and final tier, the greedy surveys.

The greedy surveys use a similar Markov Decision Process as in Tier 2, but rather than selecting large areas of sky to observe, the survey selects a single pointing at a time. No attempt is made to observe greedy scheduled observations in pairs. Since this tier is primarily used in twilight time, it only schedules observations in the redder filters r, i, z, and y.

As with the Blob tier, the sky tessellation orientation is randomized each night so the final survey is spatially dithered.

2.3. *Filter Mounting Schedule*

XXX–discuss how we decide which filters to have mounted at a given time.

3. BASIC SURVEY REQUIREMENTS

Basic survey strategy starting point and why - in more depth? Discuss metrics related to these requirements.

Probably should show that all survey strategies evaluated do / need to meet these requirements (but maybe later?)

XXX–Relevant SRD requirements. 825 observations over 18,000 square degrees, fast revisits, and astrometry

XXX–general requirement to advance all 4 pillars of Rubin science

XXX–relevant requirement to publish a list of upcoming planned observations (1?2?) hours in advance.

4. FEEDBACK FROM WHITE PAPERS AND SAC

Broad outline of points to evaluate for survey strategy, and our approach in running the subsequent experiments (this should help make sense of what comes next)

Discuss basic types of SAC recommendations.

5. OVERVIEW OF METRICS

XXX–maybe a subset of the most important metrics?

Figure 2.

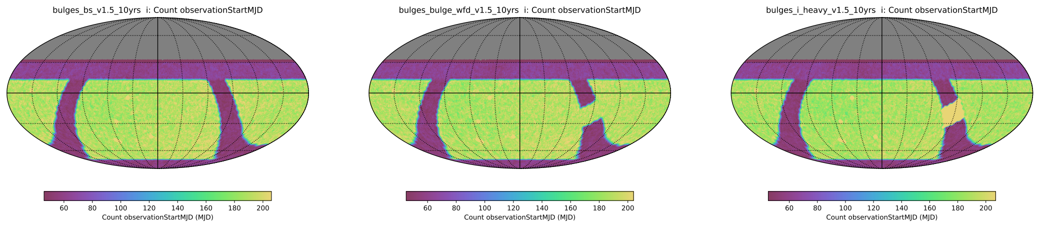

Figure 3. Series of simulations trying different bulge observing strategies.

6. INDIVIDUAL SCHEDULER EXPERIMENTS

Here we look at various experiments that explore varying a single aspect of the scheduler.

6.0.1. *Dust With Alternating*

XXX–add the alt dust plots This uses the dusty footprint, and uses a basis function to encourage the scheduler to alternate between the north and south nightly. This was originally done in the altSched (Rothchild et al. 2019). This can help keep light curve sampling optimally spaced. By using a basis function, we encourage alternating north/south, but it is not absolutely enforced, making it possible for the scheduler to avoid the moon.

There is no additional NES, however there is a strip in the north observed in g, r, i, and z.

XXX–science recap

6.1. *Bulge*

We used recommendations from the SAC for different strategies for observing the galactic bulge. These simulations use the Big Sky footprint similar to the Olsen et al white paper.

We use three footprints for bulge coverage 1) light coverage of the bulge and entire galactic plane, 2) the bulge as deep as WFD and 3) the bulge covered similarly to WFD, but with more observations in i. For each of these strategies, we run a version with natural cadence and one where we boost the priority of the bulge if it has not been observed in 2.5 days.

xxx-science recap

6.2. *DCR*

The LSST will not have an atmospheric chromatic corrector, thus difference imaging can be complicated by differential chromatic refraction (DCR). There is also potential science opportunities by being able to measure the chromatic shift in objects with sharp features in their SEDs.

These experiments look at how we could intentionally schedule a subset of images to be at high airmass so a DCR model could be built up. We test various combinations

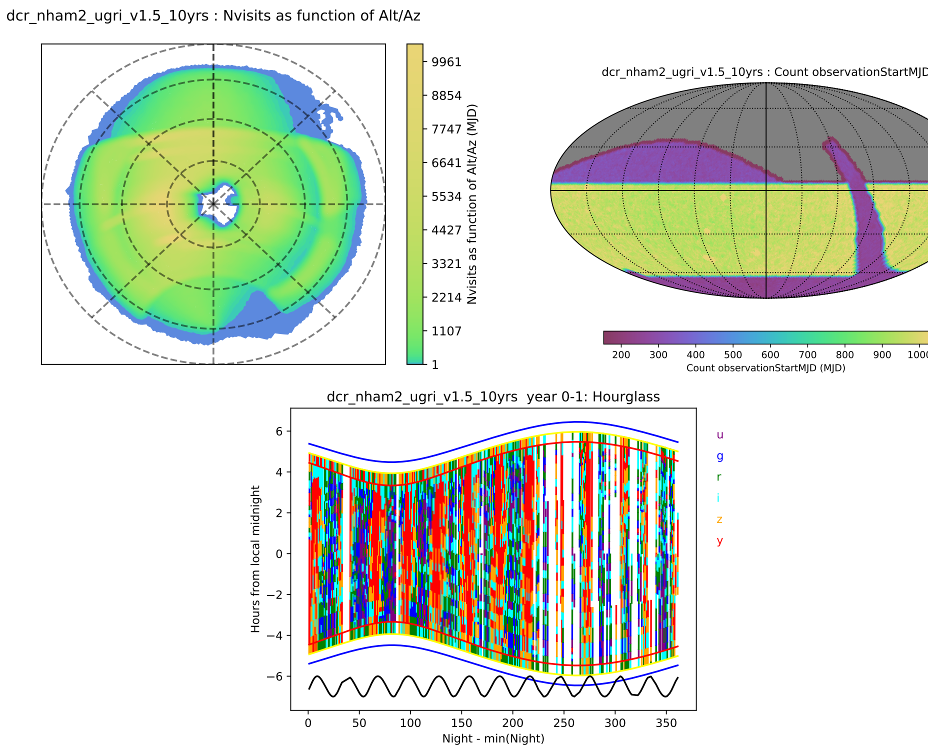

Figure 4. Intentionally taking observations at higher airmass to measure DCR.

of filters to demand DCR observations (u+g, u+g+r, and u+g+r+i), and the number of observations to take at high airmass per year (1 or 2).

Note, even with 2 high airmass observations per year, we would still expect some area of the sky to fall in chip and raft gaps. It is also worth noting that in our baseline simulation, we observe a spot on the sky in u typically 60 times, or 6 times per year. Taking 2 high airmass observations per year in u decreases the final coadded depth by 0.15 mags.

XXX–science recap

6.3. *Deep Drilling Fields*

We have run a variety of DDF strategies. Figure 5 shows the same observing season of the DDF ELIASS1 with 5 different strategies.

- AGN: This strategy takes shorter DDF sequences more often. Only ∼2.5% of visits are spent on DDFs, making the final coadded depths much shallower than other strategies.
- DESC: a strategy that split the blue and red filters to different days, emphasizing a 3-day cadence
- Baseline: Our baseline strategy where 5% of observations are allocated to DDF observations.

Name	u	g	r	i	z	y
Uniform	1.00	1.00	1	1.00	1.00	1.00
Baseline	0.31	0.44	1	1.00	0.90	0.90
g heavy	0.31	1.00	1	1.00	0.90	0.90
u heavy	0.90	0.44	1	1.00	0.90	0.90
z and y heavy	0.31	0.44	1	1.00	1.50	1.50
i heavy	0.31	0.44	1	1.50	0.90	0.90
Bluer	0.50	0.60	1	1.00	0.90	0.90
Redder	0.31	0.44	1	1.10	1.10	1.10

Table 2. Variations of the filter distribution simulated.

- Daily: Similar to the baseline, but includes short DDF sequences that can execute daily so there are no long gaps between observations
- DDF Heavy: Similar to the baseline, but XXX% of visits are allocated to DDF observations

XXX–massive table(s) of DDF, run, filter, coadded depth?

XXX–science discussion. Point out we desperately need more DDF-specific metrics from the community.

6.4. *Filter Distribution*

Testing a simple WFD-only footprint, but varying the requested ratio of observations in different filters. The different filter distributions simulated are listed in Table 2.

XXX–science recap. Note that we don't have a photo-z metric, which seems very relevant.

6.5. *Footprints*

We test a wide variation of possible survey footprints. Some of these are more realistic than others.

6.6. *Good Seeing*

These test the ability to ensure the entire WFD area is imaged in "good seeing" conditions every year, here defined as FWHM of 0.7 arcseconds or better.

These runs work well and it seems to add no particular overhead to the observing. It might make it more challenging to implement in operations, simply because the baseline simulation can simulate an entire night and pass off the list to be observed. If we want to run with the goal of collecting good seeing images, we will need to update the observing queue every time the seeing conditions change significantly, which could result in changing the upcomming observations more often than is desired.

6.7. *Short Exposures*

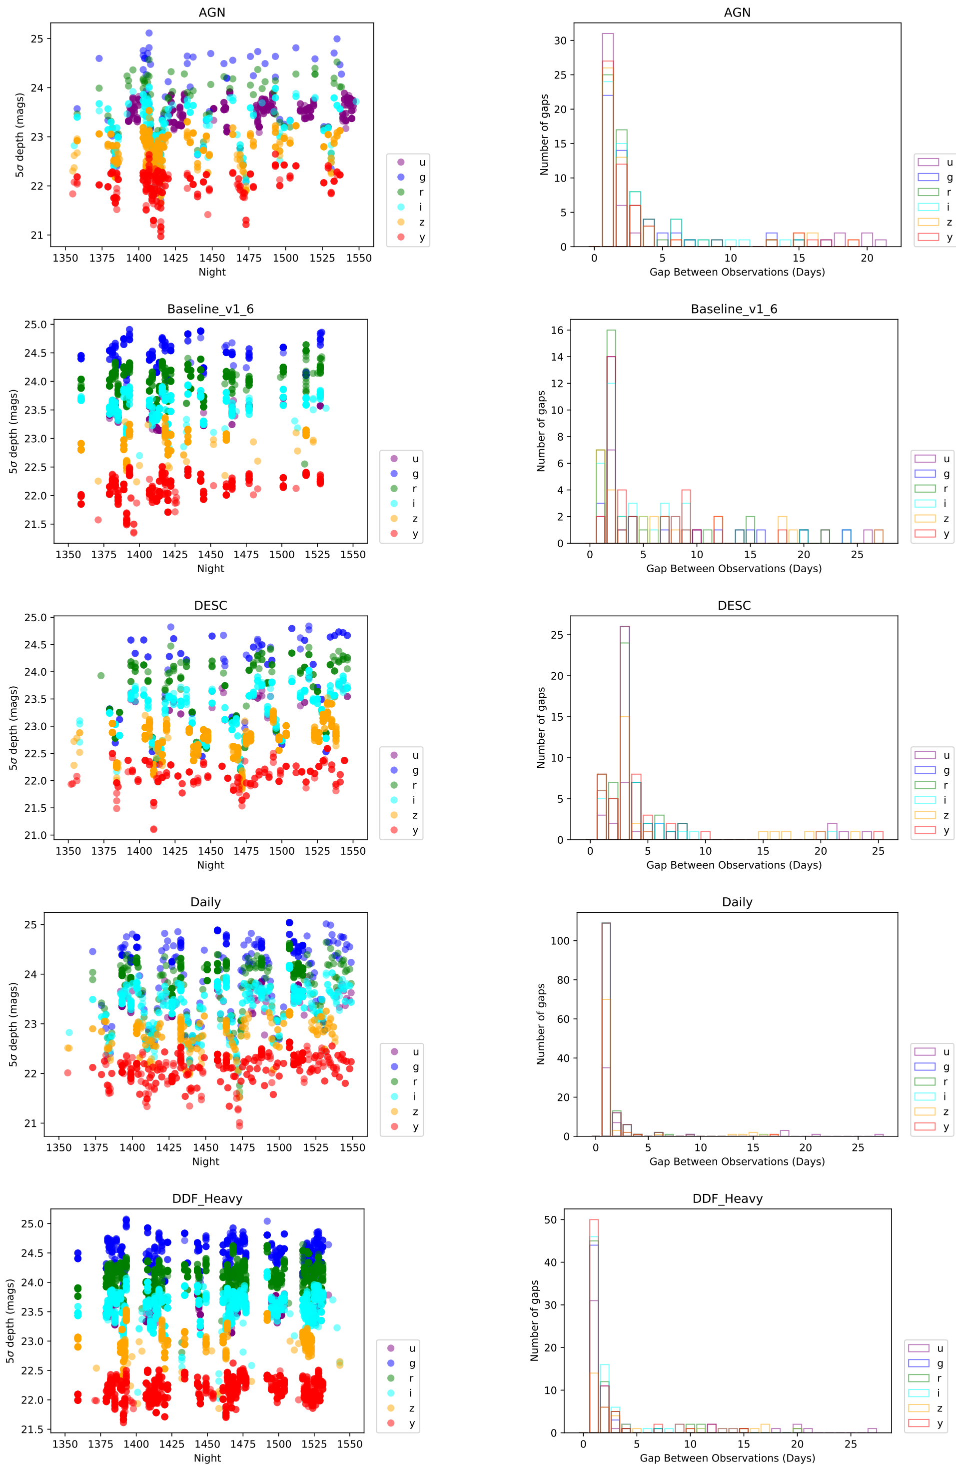

Figure 5. One observing season of the DDF ELIASS1 from 5 different DDF strategies.

We try taking additional short exposures (1s or 5s) twice or five times per year. Taking shorter exposures is a less efficient observing mode, but it seems to have little impact on the overall open shutter fraction.

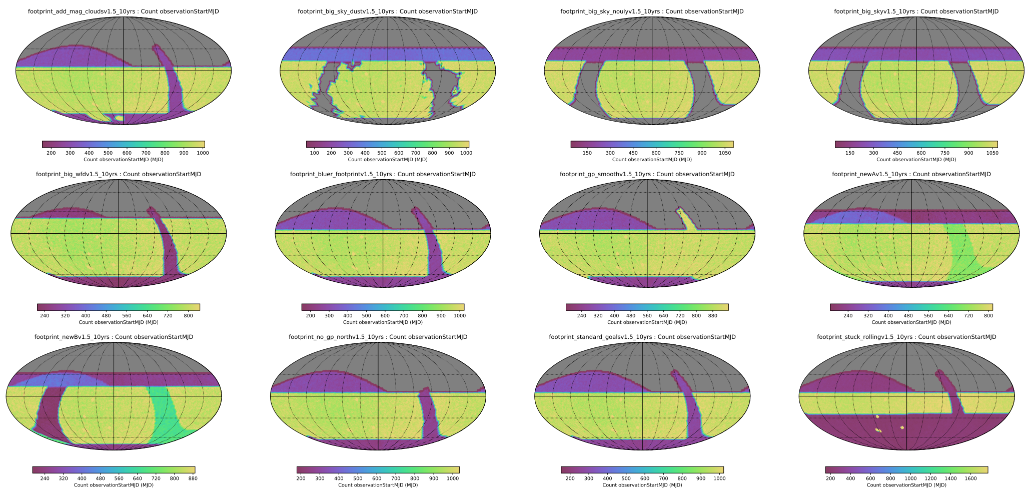

Figure 6. The different survey footprints simulated.

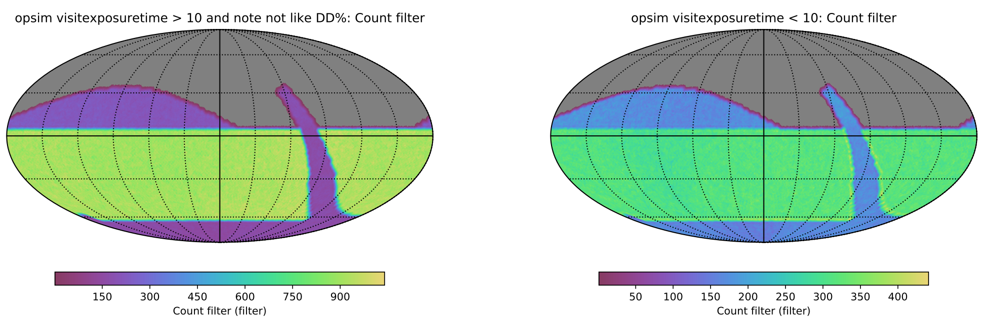

Figure 7. Results from including 5s exposures (up to 5 per year). The left shows the number of regular 30s visits (excluding DDF observations) and the right shows the number of 5s visits.

6.8. *Spiders*

We look at keeping diffraction spikes aligned along CCD rows and columns. This may result in the camera rotator angle being much less randomized than our baseline rotational dithering strategy.

xxx–Science recap: Should be almost identical to baseline.

6.9. *Third Observation*

For early identification of transients, it can be helpful to have more than two observations in a night. In these observations, we dedicate between 15 and 120 minutes at the end of the night to attempting to observe areas of sky that already have been observed.

6.10. *Twilight NEO Survey*

This is an implementation of white paper XXX, where we use twilight time to take short exposures along the ecliptic to search for NEOs.

If we dedicate all twilight time to NEO searches, we fail to meet the SRD requirements. Thus we also check running the NEO survey every 2, 3, or 4 days.

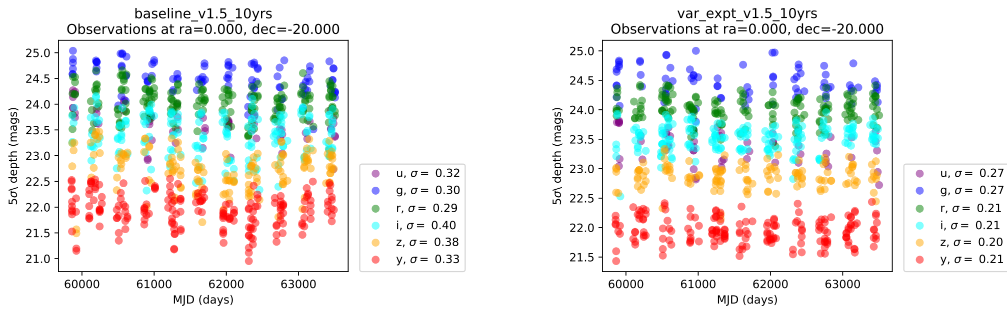

Figure 8. Comparison of a sample WFD point in the baseline and when we vary the exposure time. The individual observations depths become more uniform, especially in the redder filters that can be observed in bright time and twilight.

6.11. *Longer u Exposure Time*

The u-band observations are often expected be readnoise limited. We test doubling the u-band exposure time and cutting the number of exposures in half. This results in the u-band final coadded depth reaching ∼0.20 mags deeper. The g-band is also 0.10 mags deeper, with the rest of the filters essentially unchanged in final depth.

Note, we assume that 1x60s visit counts as 2 30s visits for the purpose of meeting the SRD value of 825 visits in the WFD area. Adopting longer exposures in u seems like a good idea, but the SRD will probably need to be modified to ensure it is not ambiguous.

6.12. *Variable Exposure Times*

We vary the exposure time based on the current conditions so individual exposures have similar depths. There is an argument that taking a full 30s visit in ideal dark time conditions results in "wasted depth", as more objects and transients will be detected, but then it will be impossible to identify them as later visits are unlikely to be as deep. Similarly, taking a 30s visit in poor conditions will result in a shallow image which will be of limited use. In good conditions, the expsoure time is allowed to shrink to 20s, and in poor conditions it can extend to 100s.

As with doing 60s u band exposures, this may require modifying the detailed specifics of the SRD as longer exposures may need to count as multiple visits.

Having variable exposure time introduces at least 8 new free parameters to the scheduler (the target individual depth for each filter), as well as the shortest and longest acceptable exposure times. As with 6.6, this would be more complicated to run in operations as the scheduler would need current conditions to calculate the modified exposure times, although the predicted sky brightness may be accurate enough.

XXX–science recap:

6.13. *WFD Depth*

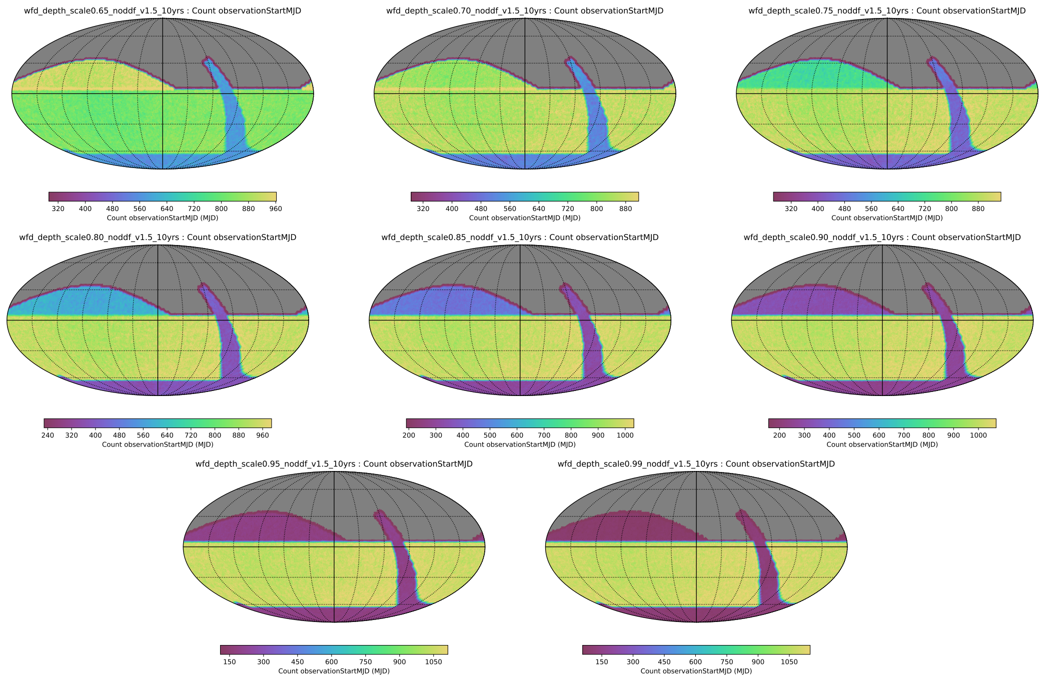

Figure 9. Varying the amount of time dedicated to the WFD region between 65% and 99% of the visits.

We vary what fraction of the observing time is dedicated to the WFD area, from 60% to 99% with and without the standard DDF surveys. Unsurprisingly, the SRD is not met if the WFD is only given 60%.

6.14. *Rolling Cadences*

6.15. *Even Filters*

The baseline simulation is fairly aggressive in switching to redder filters in bright time. This can create long gaps in light curves with no bluer observations. We have run a simulation where only the u, and g filters avoid bright time, and a simulation where only u avoids bright time. Figure **??** shows the resulting filter distributions in year one. Unlike the baseline simulations, there are no longer sections of several days where only y is observed.

XXX–science recap. Should be less depth, but maybe better long transient light curves.

6.16. *Aliasing*

There was concern that if observations were too uniformly placed on the meridian, periodic sources would be aliased. Figure **??** shows the FFT of observations at a sample WFD point in the baseline simulation. There is some aliasing at ~ 1 day which is inevitable for any ground-based telescope. The aliasing is much lower than the minion_1016 simulation that was analyzed in the Bell et al. cadence white paper.

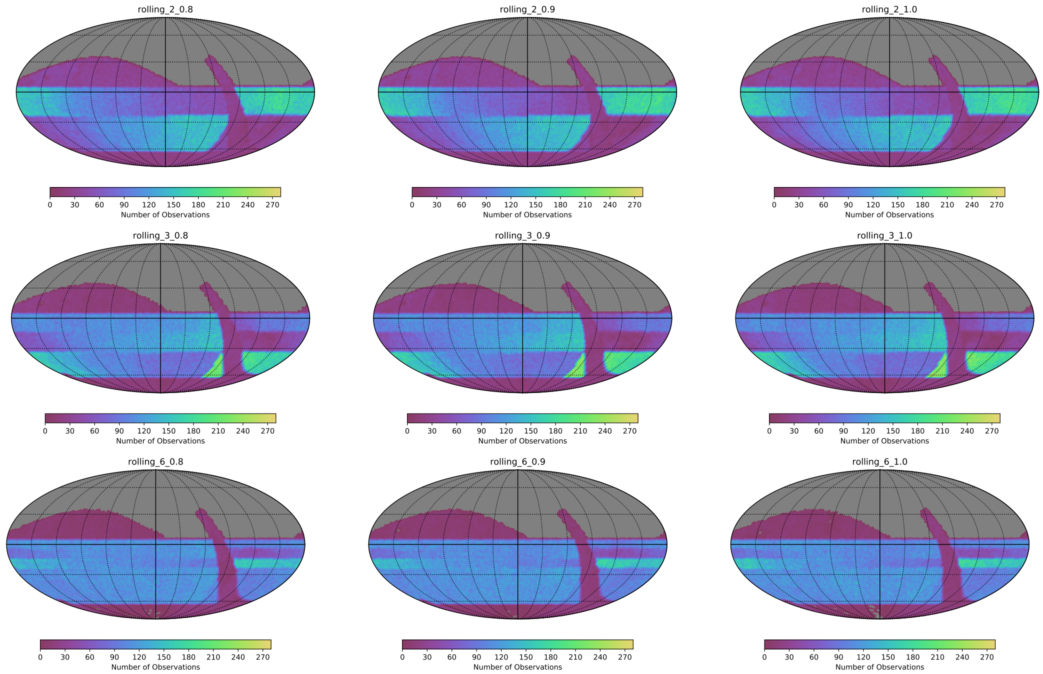

Figure 10. Rolling cadence simulations with 2 (top), 3 (middle), and 6 (bottom) rolling stripes. Here we show the observations taken from 3.5-4.5 years in the survey, excluding the DDF observations.

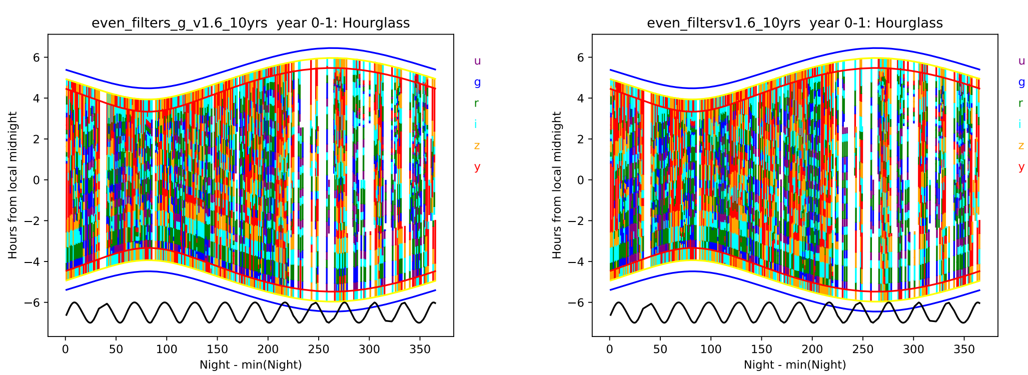

Figure 11. The filter distribution for the even filter simulations. Unlike the baseline simulations, bluer filters are observed in bright time.

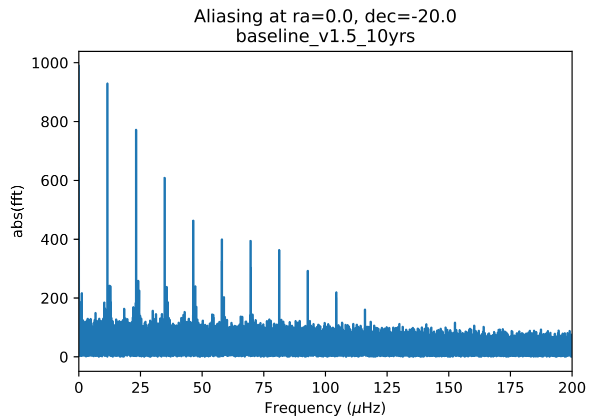

Figure 12. Aliasing at a sample position in a baseline simulation. There are peaks at harmonics of 24 hours, but this is inevitable with a ground-based telescope. The aliasing seems much lower than earlier version of OpSim where harmonic peaks could be seen past 200 μHz.

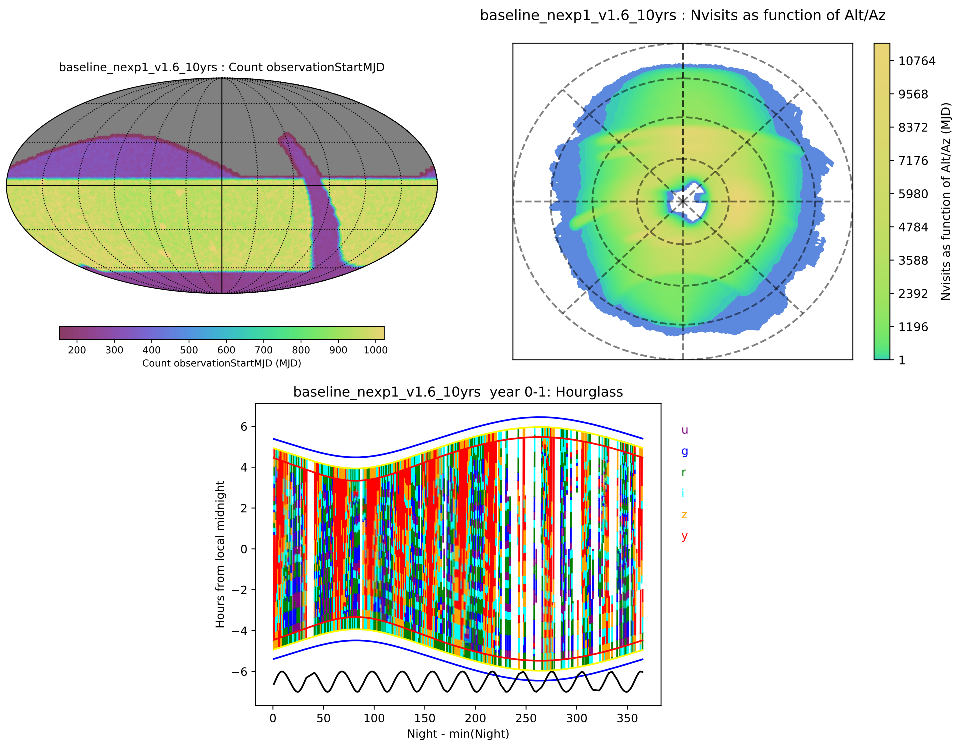

Figure 13. The baseline v1.6 simulation. The top panels show the distribution of visits (all filters) in RA/dec and Alt/Az. The bottom panel shows the first year of observations color-coded by what filter was loaded. White regions represent scheduled and unscheduled downtime as well as weather downtime. The black curve on the bottom shows the moon phase.

7. FBS RELEASE V1.6

Here we describe the runs done as part of the FBS 1.6 release. Unlike our previous simulation releases, we are trying a limited number of simulations that combine various options from previous experiments.

7.1. *Baseline*

For the baseline strategy, we set the footprint to have 18,000 square degrees dedicated the the WFD survey. The WFD has a filter distribution of u:g:r:i:z:y of 0.31:0.44:1.0:1.0:0.9:0.9. We include coverage of the Galactic Plane (GP) and South Celestial Pole (SCP). These areas are set to have 20% the number of counts of the WFD (if a spot in the WFD has 900 visits, points in the GP and SCP will have 180 visits). The GP and SCP are set to have equal number of visits in all filters. The North Ecliptic Spur (NES) is observed with only the g, r, i, and z filters. The NES area is set to have one-third the number of visits of the WFD. The filter distribution is set to g:r:i:z of 0.2:0.46:0.46:0.4.

The total breakdown of target observing time is 85% for WFD, 6% for the NES, 6% for the GP and NES, and 5% for DDFs.

While the different survey areas are covered to different depths, the baseline scheduler treats them identically and only tries to maintain the proper ratios of area coverage. This means blocks of observations can be scheduled that cover the different regions seamlessly. It also means we have no additional constraints on how the regions are observed. For example, we currently do not reserve "good seeing" time for the WFD area.

The baseline survey includes the 4 announced Deep Drilling Fields as well as a pair of fields that overlap the Euclid Deep Field South. Each individual DDF is set to take a maximum of 1% of the total visits (the Euclid pair of fields are set to a maximum of 1% combined). The standard DDF sequence is ux8, gx20, rx10, ix20, zx26, and yx20, all with 30s exposures. For any given sequence, only the five filters loaded in the camera are executed. By default, we remove the u filter when the moon is more than 40% illuminated at the start of the night.

We run 2 baseline simulations, one with 1x30s visits and one with 2x15s visits. The main difference is the additional readout time in the 2x15 version drops the open shutter fraction from 77% to 72%. This puts the 2x15s simulation close to failing the SRD FO metric, with some parts of the WFD region only reaching 824 observations (the median is still 892).

For the rest of the simulations in v1.6 we use 1x30s visits. If 2x15s visits are required there will be a significant drop in the number of visits, and areas outside of the WFD may need to be scaled back to still meet SRD requirements.

When it is non-twilight time and we are not observing DDFs, we use a Markov Decision Process to dynamically build a queue of observations. Observations are planned in 44 minute blocks (22 minutes for an initial area, 22 minutes to repeat the area). The size of the blocks can scale slightly to try and fill time before twilight (e.g., it will expand to a pair gap of 25 minutes if there are 50 minutes until morning twilight begins). All observations are taken in pairs, with potential combinations of u+g, u+r, g+r, r+i, i+z, z+y, or y+y. The ordering of the filter pairs can change depending on what filter is currently loaded (e.g., if the scheduler decides to observe a g+r sequence, the r observations will be taken first to eliminate a filter change if possible.)

The camera rotator angle (relative to the telescope) is randomly set each night between -80 and 80 degrees. The angle is set when the block is scheduled, so there can be a few degrees of drift between when the rotator angle is computed and when the observation is actually taken.

The MDP uses basis functions based on

- The 5-sigma depth (for both filters in the pair being taken)
- The footprint uniformity (again, in both filters)
- The slewtime
- Staying in the current filter
- Rewards taking 3 observations per year per filter over the entire survey footprint

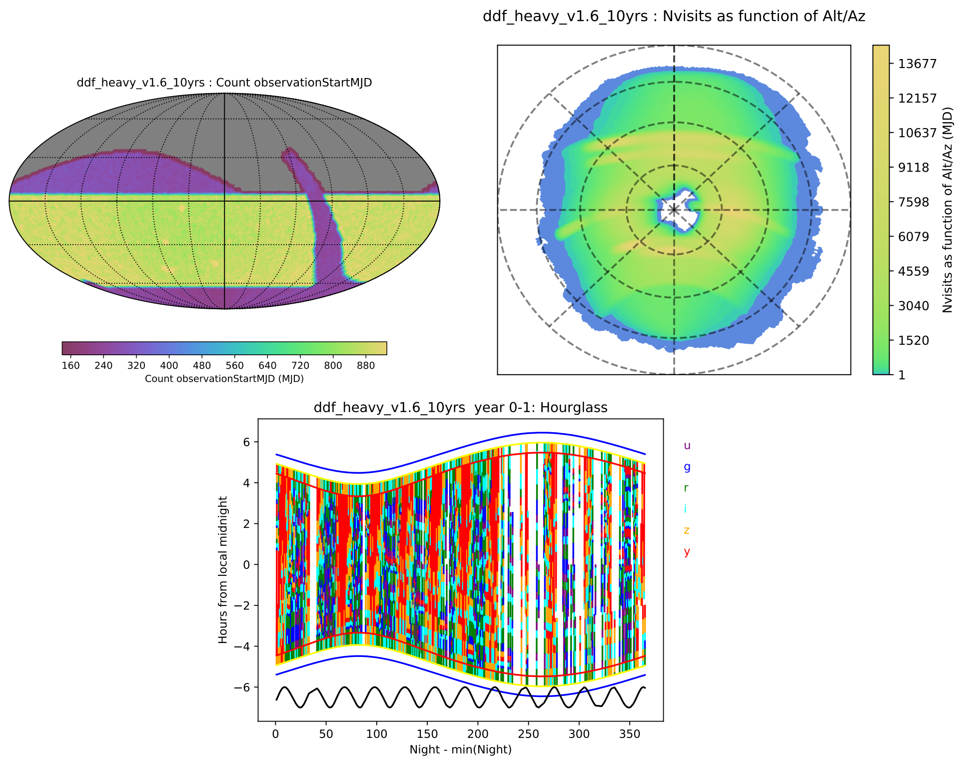

Figure 14. DDF Heavy simulation. Nearly identical to the baseline, but giving as much time as possible to DDF observations.

The MDP also includes basis functions that are simple masks

- Zenith is masked (to avoid long azimuth slews)
- 30 degrees around the moon is masked
- The bright planets (Venus, Mars, and Jupiter) are masked with a 3.5 degree radius

If the sun is higher than -18 degrees altitude, or there is not enough time remaining to take observations in pairs, the scheduler reverts to a greedy algorithm and selects observations one at a time. We use a similar MDP for these greedy twilight observation decisions.

XXX–science summary of the baseline. Compared to many of the other simulations, the baseline spends a lot of time observing the WFD, with a median of 948 visits. The higher number of visits means a faster cadence and better sampled lightcurves for objects with durations comparable to a season length.

7.2. *DDF Heavy*

This run is nearly identical to the baseline, but gives a large fraction of time to the deep drilling fields. Each of the five DDFs takes between 2.4 and 2.9% of the survey, with 13.4% of all visits being used for DDF observations. The baseline has 4.6% of visits used for DDFs. This is enough time that the WFD area near the DDFs fails to

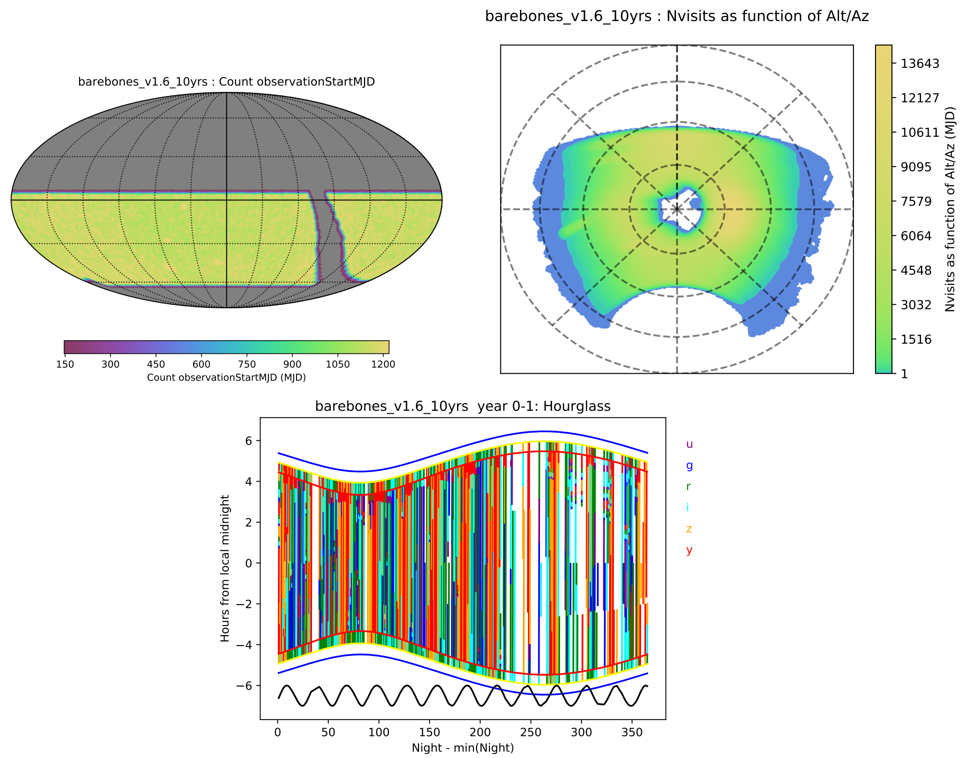

Figure 15. The barebones simulation covering just the WFD area as efficiently and deeply as possible.

reach 825 visits over 10 years, but the SRD requirement is formally still met because the median WFD point is observed 875 times.

XXX–For each of these 1.6 runs, maybe an include a science impact recap? Maybe a radar plt relative to baseline (and same scale across all of them), and a few lines of explanation of what we think happened, what metrics we still need (e.g., here we could say we need AGN and other DDF relevant metrics).

7.3. *Barebones*

The barebones simulation is an not a viable survey strategy, but provides an extreme example where we focus exclusively on meeting the SRD requirements, with little optimization for science.

The survey footprint is restricted to the baseline 18,000 square degree WFD area only. Deep drilling fields are included, but capped at $\sim 2.5\%$ of the total visits. Visits in u and y are unpaired, while the rest of the filters are paired in the same filter. This results in very few filter changes in a night.

There are a wide number of reasons why this would be a terrible survey strategy–detected transients would have no color information, photometric uber-calibration could be difficult with the galactic plane gap, a lack of solar system object because the NES is not included, etc. The main purpose is to show the scheduler can run very near the theoretical maximum for open shutter fraction, with this run reaching 80%. Also, we can note the fONv metric reaches 1,148 which is 40% higher than the

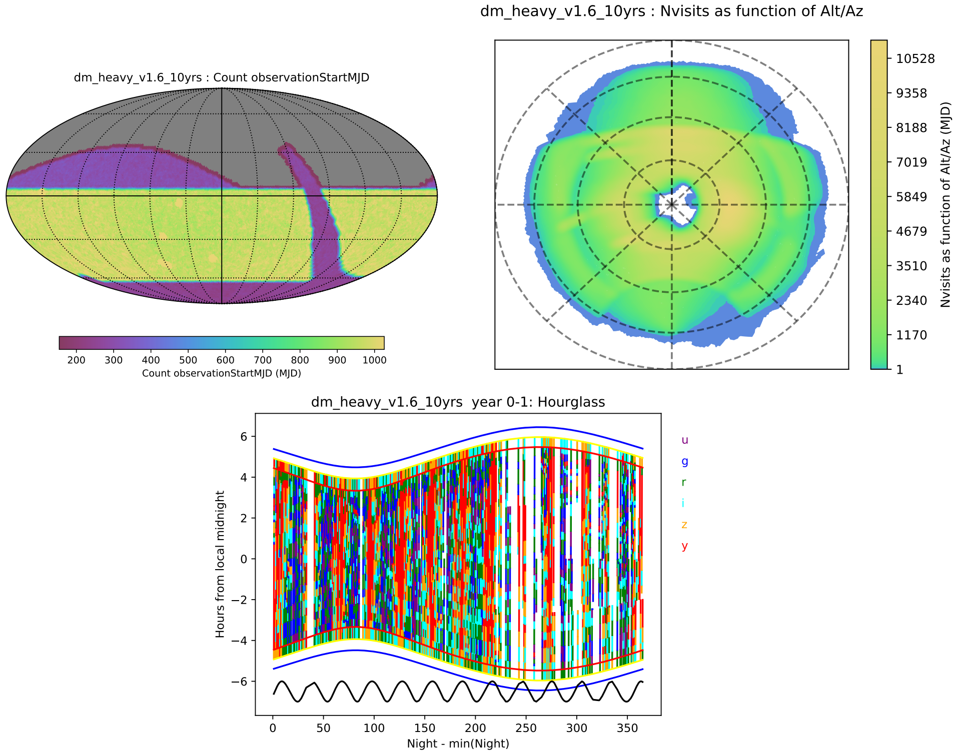

Figure 16. The DM heavy simulation. Similar to the baseline, but the alt/az plot shows how some observations are being taken at high airmass to support DCR modeling.

SRD requirement of 825. This also implies that we can observe a maximum of ~ 115 WFD visits per year in the event we want to adjust the scheduler to attempt to catch up on the WFD progress.

XXX–science wise. I bet some SNe love it. Solar System might like some of it, but the missing NES should be a killer.

7.4. *Data Management Heavy*

This is simulations includes various modifications that may be helpful for Data Management purposes. For the WFD region in u, g, and r a few images per year are taken at high airmass so that DCR correction models can be made.

The camera rotator angle is set so that diffraction spikes fall along CCD rows and columns. This helps with difference imaging so the maximum possible area can be used, but may result in weak lensing systematics.

Each year, the scheduler prioritizes taking g,r,i images of the whole sky in good seeing conditions (defined as $0.7''$effective FWHM or better).

The DDF fields use larger dithers, up to 1.5 degrees, compared to the default 0.7 degree maximum.

XXX–science summary.

7.5. *Rolling Extragalactic*

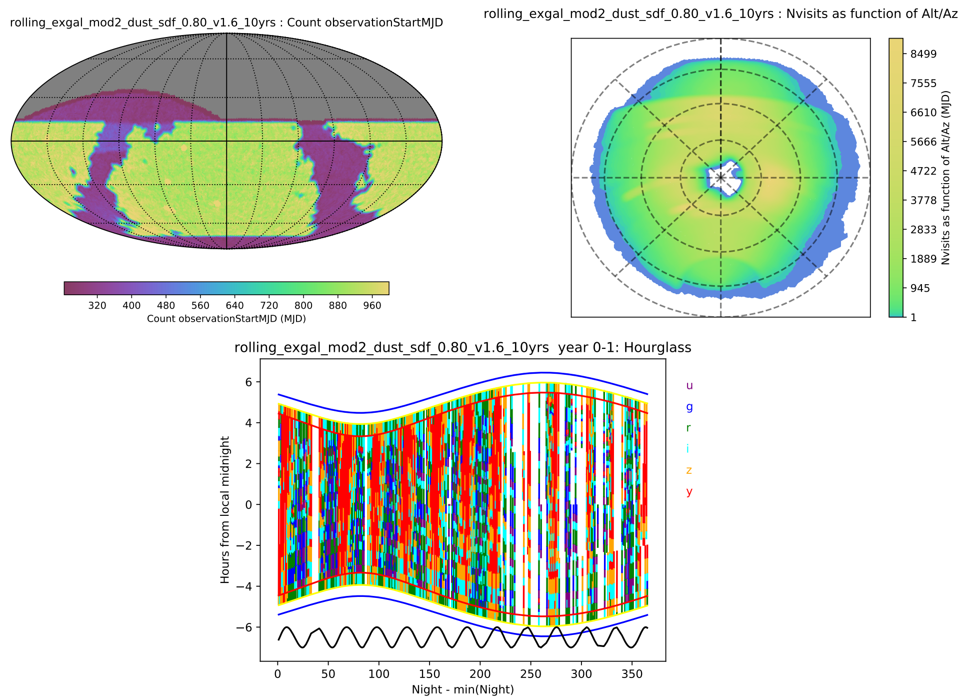

Figure 17. The Rolling Exgal simulation. The WFD area is set to be 18,000 square degrees of low extinction area.

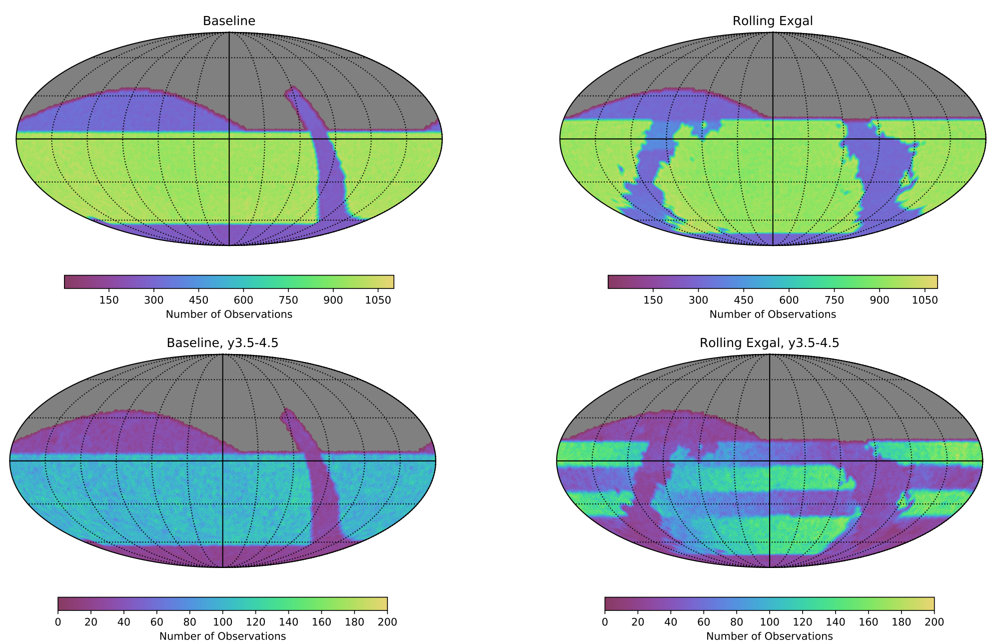

Figure 18. Illustration of rolling cadence. The top panels show the number of observations after 10 years (all filters) for the Baseline and Rolling Exgal simulations (excluding DDF observations). Both simulations have very smooth WFD coverage, with ∼900 observations. The lower panels show the number of observations taken between 3.5 and 4.5 years. The baseline WFD remains smooth, while the Rolling Exgal simulation has declination stripes of high and low counts.

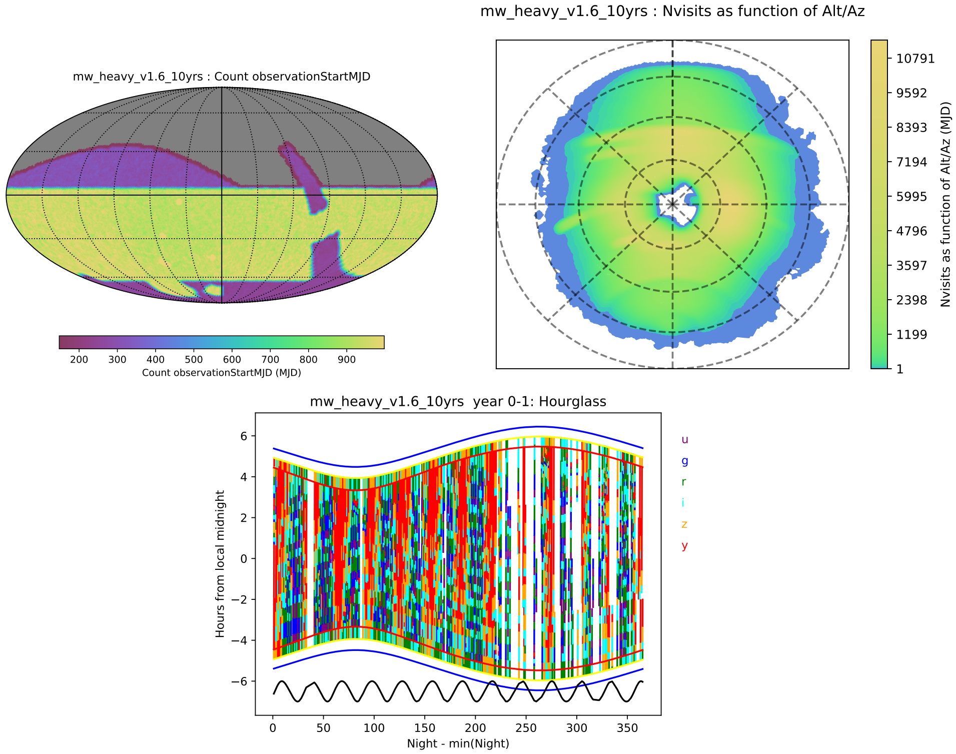

Figure 19. The Milky Way heavy simulation. Similar to the Baseline, but the bulge and Magellanic Clouds are added to the WFD area.

The rolling extragalactic is motivated by cosmological drivers. The footprint is modified so the 18,000 square degrees of the WFD are placed in low-extinction regions. The simulation also executes a half-sky rolling scheme, which should result in better sampled lightcurves for extragalactic transients.

This simulation divides the sky into quarters, and has one northern stripe and one southerns stripe with a rolling emphasis at a time. This could be preferable to a simple two-bad rolling scheme, because with the quarters a region of emphasis will always be available to northern telescopes. If we rolled with an emphasis purely on the southern half of the WFD region, ∼80% of the Rubin alert stream would become unavailable to northern hemisphere observatories for that season.

XXX–science recap.

7.6. *Milky Way Heavy*

The Milky Way heavy simulation covers the Galactic bulge, LMC, and SMC as part of the WFD area.

There is very little change in the median coadded depths compared with the baseline since the extra WFD area is added to a region of the sky that is under-subscribed in the baseline. In the baseline simulation, there are an excess of observations near the galactic plane, so covering the bulge is "free", in the sense that it uses these excess pointings to cover the bulge.

XXX–science recap. Should be a huge boost in microlensing events.

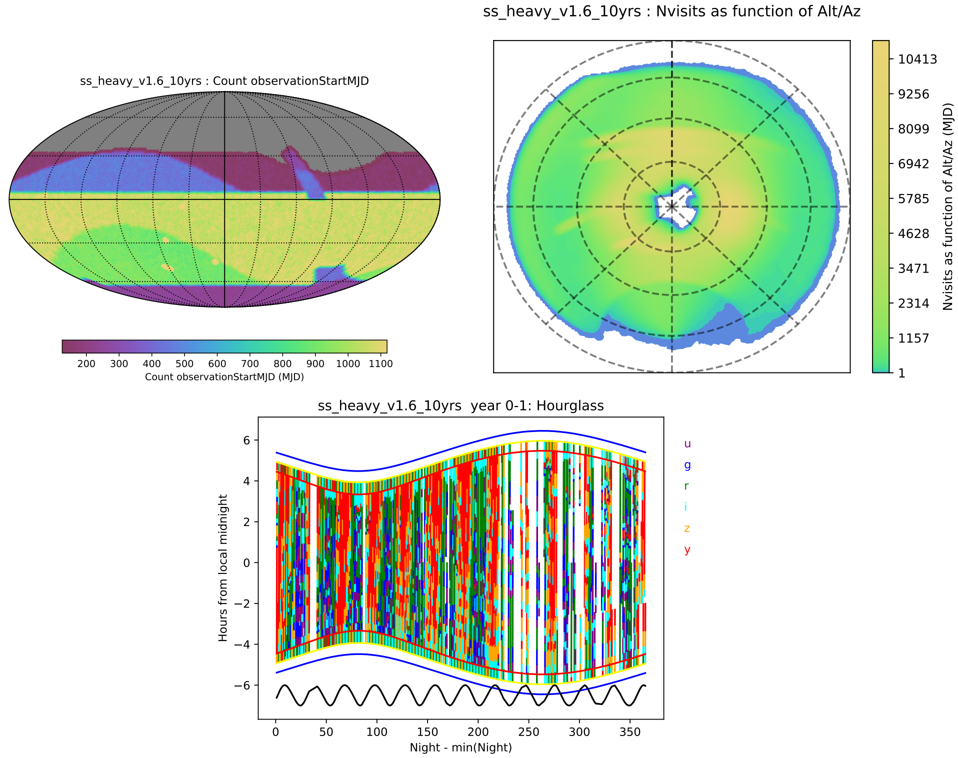

Figure 20. The Solar System heavy simulation. The high airmass observations are twilight NEO observations.

7.7. *Solar System Heavy*

For the Solar System Heavy simulation, the footprint in modified to include ecliptic plane coverage through the galactic plane.

A fraction of twilight time is used for a NEO survey. The NEO survey uses very short exposures at high airmass. Note, a NEO survey taking short exposures will drastically increase the data throughput of the system. DM needs to check if this mode of observing would be feasible. We also need to check with the camera team that taking short exposures for an extended time will not be a thermal issue.

This simulation only uses i,z,y in twilight time, making sure we observe more r-band in non-twilight and in pairs. It also includes r+r pairs in non-twilight time. For regular 1x30s visit twilight observations, we avoid observing the ecliptic, thereby ensuring they are always taken in pairs in non-twilight time.

XXX–science recap. Hopefully this helps SS science!

7.8. *Combo Dust*

This simulation attempts to improve several science cases compared to the baseline simultaneously. The footprint used here starts with defining the WFD area as 18,000 square degrees with low extinction. Then an additional 2,000 square degrees are added to WFD to cover the bulge, the ecliptic through the galactic plane, the LMC and SMC, and an outer Galactic plane region. Dusty areas of the sky and the South Celestial Pole are covered at about one-quarter the WFD depth. The NES is covered

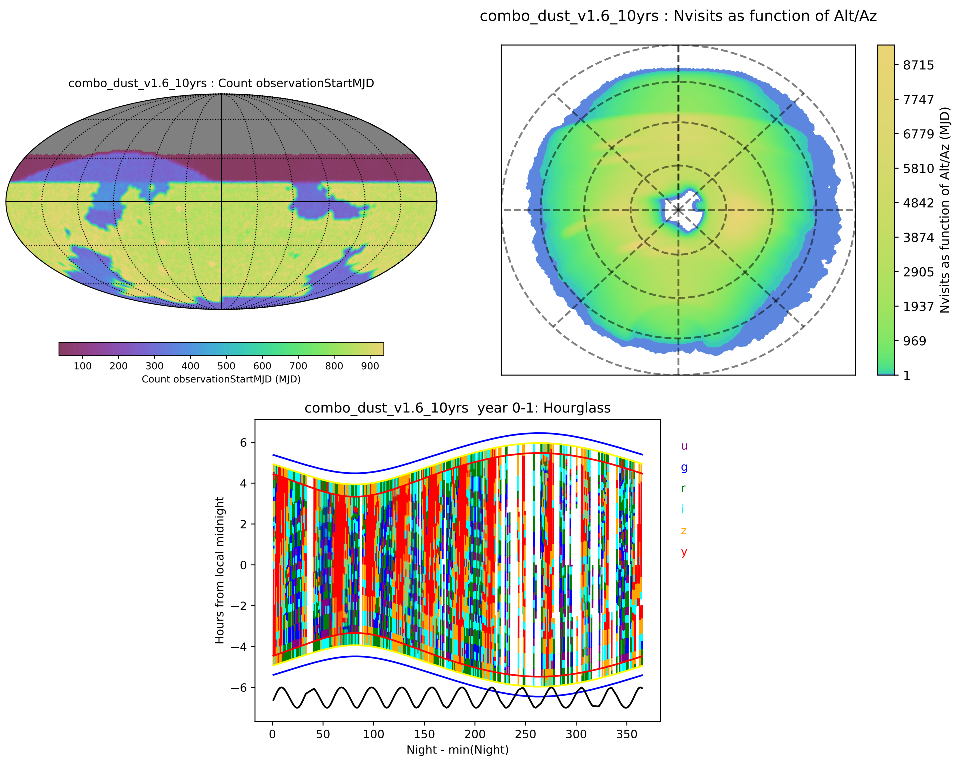

Figure 21. The Combo Dust simulation. Similar to the Rolling Exgal simulation, but the WFD is expanded to include the bulge and ecliptic, Magellanic Clouds, and an anti-center bridge.

filter	Baseline	Baseline 2 snaps	Barebones	DDF Heavy	DM Heavy	MW Heavy	Rolling Exgal	SS Heavy	Combo Dust
	(mags)				$m_{\rm Baseline} - m_{\rm Sim}$				
u	25.86	0.24	-0.13	0.08	0.11	0.02	0.11	-0.02	0.12
g	26.97	0.11	-0.15	0.09	0.12	0.01	0.10	0.07	0.14
r	26.95	0.08	-0.12	0.08	0.07	0.01	0.10	0.05	0.14
i	26.40	0.07	-0.17	0.11	-0.01	0.01	0.11	0.11	0.15
z	25.67	0.06	-0.12	0.08	-0.01	0.01	0.11	0.02	0.11
y	24.90	0.06	-0.14	0.06	0.04	0.01	0.09	0.03	0.09

Table 3. Difference in median depths compared to Baseline for v1.6 simulations. Negative values indicate deeper depths.

in g, r, i, and z. The footprint also includes very light coverage to the northern limit of the telescope in g, r, and i so there can be templates for ToO events on the entire accessible sky. This simulation includes the same half-footprint rolling scheme as Rolling Extragalactic.

The footprint has 35 free parameters for setting region locations and filter ratios. Many of these have have been set by eye or use historical values of questionable providence.

XXX–science recap:

8. INDIVIDUAL VISIT LENGTH

What to do - 1x30s vs. 2x15s? 1x30s much more efficient (show rough calculation of overhead) than 2x15s, but may have drawbacks due to cosmic ray rejection and potential to miss very rapid transients (or WD detection .. ref white paper). Subtle drawback that 2x15s gives the same ”midpoint exposure time” across FOV, 1x30s does not.

Show difference in 1x30s vs. 2x15s in whatever is our ’standard baseline’ at this point.

There has been thought of using a variety of exposure times if we use two snaps (e.g., 5s + 25s). Because there are not plans to release catalogs from individual snaps, it’s not clear if this would enable much new science.

Show effect of 7% loss in efficiency when attempting to combine minisurveys in various configurations (assume we will find some combinations possible with single exposure visits that are impossible with two snaps).

Also possible to use variable exposure time depending on seeing and sky brightness conditions. Shorter exposures in good conditions keeps us from observing “wasted” depth, letting us take longer exposures in poor conditions. This does introduce a host of new free parameters (an ideal target depth for each filter and minimum and maximum exposure times). This would might require rewording the SRD to ensure, e.g., that 20s visits in good conditions count for the number of visit requirement.

Relevant metrics: total number of visits, number of visits per field/filter

9. INTRA-NIGHT CADENCE

What to do for visit sequence within a night? White paper support for multiple filters within a night (except TNOs maybe?). Potential drawbacks - less efficient (show effect on efficiency). This applies to WFD primarily, but we’ve applied to any survey that did not have their own specifications (so, everywhere).

Extension of pairs to u band and y band (show effect).

Relevant metrics: inter-night visit gaps and SN discovery, SSO discovery/characterization, transient and variable discovery (??), number of visits

10. WIDE-FAST-DEEP FOOTPRINT

What to do for WFD footprint? SRD not specific, DESC want low-extinction sky (and depth), but WFD is generally the area of sky that receives the most visits, so generally other science will also benefit from more visits to their relevant areas (particularly galactic plane .. for time-domain studies primarily, not depth)

Relevant metrics: area of sky with 825 visits (under particular restrictions, like total coadded depth and individual image seeing and dust extinction), number of galaxies, number of resolved galaxies, SSO discovery, transient and variable star discovery, astrometry in the galactic plane (?)

11. ROLLING CADENCE

Motivation for a rolling cadence (more frequent visits in some years)

Different options for rolling and explanation of how implemented

Should really include discussion of recovery from bad weather years and simulation of same

Relevant metrics: Maintain astrometry requirements, SN discovery, SSO discovery and characterization, Transient and variable discovery, uniformity of coadded depth / number of visits,

12. NORTHERN MINISURVEYS

Add extension to cover Euclid/DESI with various numbers of visits

Observing NES

Effect of adding or removing these minisurveys

Relevant metrics: SSO discovery and characterization (particularly active asteroids), depth and number of visits through remainder of North

13. SOUTHERN MINISURVEYS

Add extension over south celestial pole, LMC/SMC with various numbers of visits

Effect of adding or removing these minisurveys

Relevant metrics: number of visits and coadded depth over SCP, discovery of variables in LMC/SMC (see Olsen white paper for metrics?)

14. LOW GALACTIC LATITUDES

Discussion of definitions from SAC and recommendations for visits

Effect of adding or removing these minisurveys

Relevant metrics: number of visits, astrometry in bulge, discovery of variables/transients/microlensing in bulge (?)

15. TWILIGHT OBSERVING

Discuss need for twilight observing to meet SRD goals (weather, total amount of time available)

Add NEO twilight survey, add DCR white paper (season extension visits?)

Effect of adding or removing these minisurveys

Relevant metrics: NEO discovery, number of visits and coadded depth (and uniformity) in WFD, measurement of DCR, season length

16. DEEP DRILLING FIELDS

Discuss purpose and how these are scheduled (very different from other fields)

Discuss potential cadences (AGN/ DESC) and how these differ, and our combination of the two

Discuss timing issues with oversubscription (and how much of a problem this could be, what if worse weather?) – include location of fifth DD field

Effect of adding or removing these minisurveys

Relevant metrics: number of visits and coadded depth for DD, SN detection in DDFs, AGN detection in DDFs *[solar system minisurvey DDF?]

17. TOO MODES

Discuss impact of ToO, and how we could implement ToOs in scheduler (various modes: straight to queue by hand or set up known program and supply trigger, etc. – that we're evaluating the second?)

Any ToO survey should also take into account that chip and raft gaps mean full sky coverage will require multiple images with spatial dithering.

Discuss how we can have a low coverage region to the north to maintain templates for all possible ToOs, or we could decide ot only search for ToOs that are likely to be in the WFD area.

Relevant metrics: frequency of achieving ToO observations, number of visits and coadded depth in other surveys (WFD or other minisurveys that may be in particular contention)

18. FURTHER OPTIMIZING

Somewhere in here we probably ought to talk about optimizing the parameters for each run, and doing bigger sweeps across parameter space. That would easily expand each of the above options by many factors.

XXX–need to optimize basis function weights for both the blobs and the greedy algo.

XXX–can also discuss pre-scheduling the DDF sequences here.

19. CONCLUSIONS

Hopefully here we pare down the evaluation of 100s of runs (like promised) to a set of between 10 to 20 (if this is possible, after combining along different axes). The results should come with some basic comments about what's particularly good or bad in each of these areas and how we arrived at these general options.

20. OUTSTANDING QUESTIONS

Here we go through some of the outstanding questions that the SCOC and scientific community can help resolve in order to converge on a final scheduler strategy for the Rubin Observatory.

20.1. *Exposure Time(s)*

We will probably need on-sky data to make a final answer to this question, but we need to eventually decide how many snaps to take in a visit. We have run the baseline simulation with both 1x30s visits and 2x15s visits. Another possibility is using variable exposure times to make the single visit depths more uniform.

Other questions related to exposure time

- Should we change the u-band to default to 60 second exposures to ensure they are not readnoise dominated? This might require decreasing the SRD 825 visit value. This choice would also severely limit u band time domain science (e.g., TDE early detection)
- Should we include some very short exposure time exposures. That would let us have better tie-in with other surveys (e.g., Gaia). It is relatively little exposure time, but the readout time means it is a low-efficiency way to operate the telescope.
- Should we decrease the exposure time in twilight to keep the saturation level reasonable?
- Should we use variable exposure times so individual exposures have more uniform depth? In poor observing conditions, we would have fewer exposures that were londer and in good conditions we would have more observations that are shorter.

20.2. *Pairs and Filter Choice*

There is a strong preference to take observations in pairs. Closely spaced observations let the pipeline identify moving objects. Similarly, observations in different filters are essential for transient classification.

The baseline survey (and most of our other experiments), take pairs in neighboring filters (e.g., u+g, g+r, etc). We should verify that this is a good pairing strategy. Similarly, we have done experiments where we attempt to observe a third observation in a night.

Taking pairs in different filters does increase time spent changing filters, but it's something like a 4% hit that seems totally worth the science gain so far.

Our baseline strategies heavily prefer to take y-band observations in bright time. While this is optimal for the possible SNR, it can result in long gaps between observations in bluer filters.

Similarly, we could expand or constrict the filters we pursue in twilight time. XXX– some 1.5 runs I think that vary which filters we run in twilight.

20.3. *Survey Contingency*

How much contingency should we aim for when designing the survey strategy? Currently, with what we believe is a conservative weather closure policy, we can meet SRD requirements with 2x15s visits, but can cover a larger footprint and do more science cases with 1x30s snaps.

20.4. *Deep Drilling Fields*

We have run a variety of Deep Drilling strategies. The DDF strategy is largely separable from the rest of the survey design, and we have a number of proposals for DDFs that we have yet to explore (e.g., rolling DDFs where a single DDF is completed in one observing season). We have started experimenting with pre-scheduling DDF observations.

- What fraction of the survey should be dedicated to the DDFs?
- Should DDFs be preferentially executed in dark time, or is it more important to maintain cadence?
- Where should the DDFs be placed (can we finalize the 5th DDF as a Euclid double-pointing)?
- What is the preferred dithering strategy (spatially and rotationally) for the DDFs? There is tension in that DM generally prefers larger dithers for calibration and co-addition purposes, while science cases prefer smaller dithers to preserve the area that reaches the deepest levels.
- Should we try "rolling" the DDFs, completing DDF observations in a field in only a few years?

20.5. *Rotational Dithering*

By default, we select a random camera rotation angle (wrt the telescope) nightly. This creates minimal additional slewtime, and seems to provide adequate angular randomization. We currently have no science metrics that depend on the angular distribution, and this should be something very important to weak lensing science (although we do not have a metric to measure this).

We have also experimented with setting the camera rotation angle to ensure stellar diffraction spikes fall preferentially along rows and columns.

- How should we rotationally dither visits?

20.6. *Spatial Dithering*

For the wide area regions we have had excellent results randomizing the tessellation orientation nightly. This does result in a small percent of time being spent observing outside the desired survey footprint. The alternative would be to limit the amount one dithers out of the footprint, but then one risks imprinting systematics on objects near the footprint border (e.g., an object is never observed in the center of the focal plane, only by outer rafts).

20.7. *Survey Footprint*

Perhaps the biggest question, what should we set the survey footprint to be?

- How should we cover the Galactic plane?
- How should we observe the Galactic bulge?
- Should we avoid areas of high dust extinction for the WFD area?
- What is the ideal filter distribution to use? It would be nice to have a photo-z metric to help make this decision.
- What is the ideal filter distribution in the GP and SCP?
- Should we cover the LMC and SMC as part of the WFD survey? As their own DDF-like survey? We have few metrics that touch on LMC/SMC science directly.
- Should we add area in the north to overlap with Euclid, WFIRST, and/or DESI?

Once the general survey footprint is decided, we can fine-tune the footprint (e.g., tapering the WFD region slightly around the RA with multiple DDFs, and flaring at under-subscribed RAs).

20.8. *Rolling Cadence*

We have gone through several iterations of rolling cadence, and now have started to converge on a technique that does not seem to impact the final survey depth.

- Should we use a rolling cadence strategy?
- Should we roll just the WFD area, or other regions as well?

20.9. *Best Use of Twilight Time*

Our baseline simulation uses twilight time to fill in WFD observations in redder filters (*rizy*). We can use some of the time to conduct a NEO survey. We can also vary which filters get used in twilight time. The baseline greedy algorithm used in twilight is known to be rather unstable, so we could also try running more contiguous blocks in twilight. We could also emphasize targeting areas that have already been observed 4 or more times in the night, potentially gathering important color information for a small number of transients.

20.10. *Target of Opportunity*

Currently, the only expected ToO use of Rubin observatory is follow up of gravitational wave detections.

- When should Rubin interrupt observations to look for GW optical counterparts?
- Do we look for GW events in the WFD area, or anywhere on the sky?
- Should we expand the survey footprint so we have image differencing templates over the entire accessible sky, in at least a few filters?
- Should Rubin plan on observing the entire light curve of ToO events, or make observations primarily for detection/classification and leave detailed follow up to other observatories?
- What filter combination and dither strategy (filling chip and raft gaps) should be used for observing ToO triggers?

20.11. *Image Differencing Templates, DCR*

Do we need to do anything special to ensure we have adequate image templates? A certain number of observations per year? A certain fraction of images taken in good seeing conditions?

If we need to start considering image quality, that makes it more difficult to simulate a night ahead of time and maintain the list of upcoming observations.

Should we intentionally extend to high airmass to facilitate DCR modeling? Note that in the baseline, we only image a location in the WFD region $\sim$9 times per year in g and $\sim$6 times in u. Also, we have chip and raft gaps, so if we want to build a DCR model for the entire sky in g, we might be dedicating 1/3 of the g observations in a year to DCR. If we switch to 60s u band exposures, there would be no observations beyond building the DCR model.

There have been claims that measuring DCR can be used for science. We do not have any metrics that demonstrate any gains, and the loss of depth is noticeable. In theory, we could combine the DCR measurements to extend the season length of observations as well (e.g., only take DCR template images near twilight in the direction of the sun).

20.12. *Satellite Megaconstellations*

Starlink is poised to launch thousands of LEO satellites. Observations so far imply that final-orbit Starlink satellites should not saturate Rubin exposures, and thus can be masked fairly easily in the image reduction pipeline.

Do we need any further satellite mitigations? Will NEO twilight surveys still be viable in the presence of megaconstellations, or should we use twilight strategies that avoid the horizon?

Figure **??** shows how illuminated megaconstellations in LEO would leave numerous streaks on Rubin images.

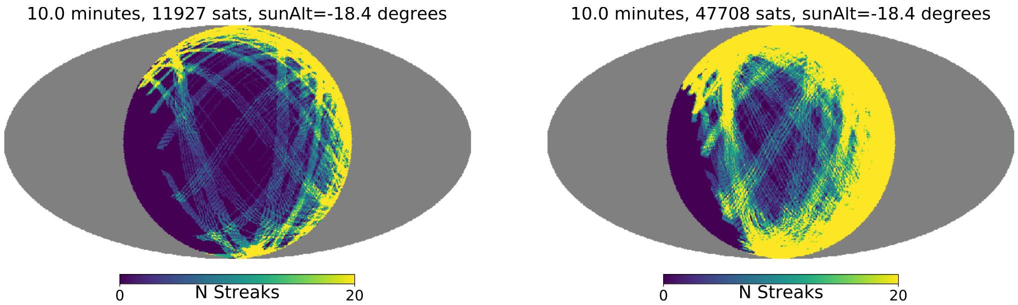

Figure 22. Alt/az projection of simulated satellite megaconstellations as seen from the Rubin Observatory site after twilight has ended.

20.13. *Aliasing*

- Are we taking observations at a large enough hour angle range that we do not need to implement further efforts to prevent aliasing of periodic sources?

20.14. *Classification of Transients*

One area where we should dedicate more work on is checking how well different survey strategies enable transient classification. This is similar to the PLAsTiCC challenge where a simulated survey was used to generate light curves and then fed to multiple classifiers. Here, rather than vary the classifiers, we vary the survey realization.

- Is the survey strategy adequate for classifying transients, or should we place more emphasis on getting more than two observations of a point in a night?

21. ACRONYMS

Acronym	Description
AGN	Active Galactic Nuclei
CCD	Charge-Coupled Device
COSEP	Community Observing Strategy Evaluation Paper
DCR	Differential Chromatic Refraction
DDF	Deep Drilling Fields
DESC	Dark Energy Science Collaboration
DESI	Dark Energy Spectroscopic Instrument
DM	Data Management
FFT	Fast Fourier Transform
FOV	Field of View (LSST FOV is 3.5 sq deg)
FWHM	Full Width at Half-Maximum
GP	Galactic Plane (galactic plane modification to survey footprint)
GW	Gravitational Wave
LMC	Large Magellanic Cloud
LSST	Legacy Survey of Space and Time (formerly Large Synoptic Survey Telescope)
MAF	Metrics Analysis Framework
NEO	Near-Earth Object
NES	North Ecliptic Spur (northern extension in ecliptic plane to survey footprint)
OpSim	Operations Simulation
RA	Right Ascension
SAC	Science Advisory Committee
SCOC	Survey Cadence Optimization Committee
SCP	Southern Celestial Pole (southern extension to survey footprint)
SMC	Small Magellanic Cloud
SN	Supernova
SRD	LSST Science Requirements; LPM-17
SS	Subsystem Scientist
SSO	Solar System Object
TDE	Tidal Disruption Event
WD	White Dwarf
WFD	Wide Fast Deep (standard 'universal' footprint)
XMM	X-ray Multi-mirror Mission (ESA; officially known as XMM-Newton)
arcsec	arcsecond second of arc (unit of angle)

REFERENCES

Naghib, E., Yoachim, P., Vanderbei, R. J., Connolly, A. J., & Jones, R. L. 2019, AJ, 157, 151

Rothchild, D., Stubbs, C., & Yoachim, P. 2019, PASP, 131, 115002