

Operation Planning for the Large-aperture Synoptic Survey Telescope

A.Saha^a, Chris Smith^b and Todd Boroson^a

^aNOAO, 950 N. Cherry Avenue, Tucson, AZ, USA

^bNOAO, Casilla 603, La Serena, CHILE

ABSTRACT

The Large-aperture Synoptic Survey Telescope will repeatedly image a large fraction of the visible sky in multiple optical passbands in a way that will sample temporal phenomena over a large range of time scales. This will enable a suite of synoptic investigations that range in temporal sampling requirements from the detection of near Earth asteroids (minutes), through discovery and followup of supernovae to long period monitoring of QSOs, AGN and LPVs (years). Additionally, the data must be obtained in a way to support programs aimed at building up deep static images of part or all of the sky.

Here we examine some of the issues involved in crafting an observing scheme that serves these goals. The problem has several parts: a) what is the optimal time sampling strategy that best serves the desired temporal range? b) how can a chosen time sampling sequence be packed into an observing scheme that accommodates all pointings and ‘whiteout’ windows (daytime, lunation period)? c) how vulnerable is such an observing plan to realistic models of disruption by poor observing conditions and weather? d) how does one build in the most economical contingency/redundancy to i) mitigate against such disruption and ii) reserve time for recovery and followup of transient phenomena (e.g. gamma-ray bursts, supernovae)?

In this article we touch upon several of these issues, and come to an understanding of some of the limitations, as well as areas in which scientific priorities and trade-offs will have to be made.

Keywords: LSST - Operations Model

1. INTRODUCTION

The Large-Aperture Synoptic Survey Telescope (LSST) is a proposed facility that will have an effective aperture greater than 6.5m and a focal plane providing a field of view of ≈ 7 square degrees. Exposures of approximately 20s duration in dark skies will reach a limiting magnitude of $R \sim 24$. The telescope is expected to take a 20s exposure and move to the next (adjacent) pointing every 30s. The desire is to repeatedly image as much of the visible sky as possible, with the primary goal of addressing time varying phenomena: discovering and cataloging moving objects, particularly near earth asteroids (NEAs) and Kuiper belt objects (KBOs); discovering transient phenomena, particularly supernovae and GRB afterglows; and discovering variable phenomena on all possible time scales. The data from the available epochs will also be “stacked” to obtain very deep sky surveys.

The eventual observing and operation strategy will emerge from the relative importance that is assigned to the very many science goals that such an experiment can address. At this time it is pertinent to develop the techniques that must be brought to bear in the planning exercises. Specifically, it is imperative that we understand the limitations and trade-off consequences in how we operate such a telescope. In this context, we present here some simple case studies in planning observing cadences.

We must recognize that the search for NEAs is likely to be a high priority for this telescope. This problem, which must chase *moving* objects over large portions of the sky, requires both frequent revisits to the same fields and coverage of a large fraction of the sky. At issue is the ability to identify the same object as it moves up to

Further author information: (Send correspondence to A. Saha)

A. Saha: E-mail: asaha@noao.edu, Telephone: 1 520 318 8288

several degrees a day. A conservative approach that has been proposed requires a given field to be imaged 9 times per month (lunation period), in sets of three exposures grouped in a single night. Another less demanding plan requires only 6 images per month, but employing more than one passband, and a cadence of 2 images on each of days 1, 2 and 10. It is important to assess what impact this will have on other time domain projects so that we can make informed trade-offs. As we shall see, the eventual choice of cadence for the NEA search has a very large impact on the efficiency of other science projects with the *LSST*.

This study examines the efficacy of various observing cadences within the boundary conditions of the projected telescope and diurnal, lunar and solar periodicities in the observing window. How sensitive are we to the number and frequency of observing epochs and to missed epochs due to weather and breakdown? Can we identify an overall operation strategy that permits good recovery from missed epochs? We proceed here in the context of detecting variability and transient behavior in “fixed” objects. The case for moving objects requires simulation experiments that are beyond the scope of this paper, but we discuss the impact of the sampling strategies for NEA search using the two models mentioned above.

In this paper we do not explicitly model the impact of weather and other unpredictable hindrances to observing. We shall see that the first level of concerns to address are prominent enough even without considering such factors. We recognize that this area must be addressed in the next round of refinements.

2. THE TEMPORAL BOUNDARY CONDITIONS

From a single site, one can only observe during a fraction of the 24 hour day. A given object is available only during a fraction of a year. The phase of the moon additionally affects and modulates observing windows. These limitations are fully predictable. In addition, there are unpredictable exigencies such as weather, observing conditions, and telescope and instrumentation breakdown. These factors set hard limits on how well we can cover a spectrum of timescales when investigating time-varying phenomena. In this section we develop 1) how the power spectrum in frequency space helps us to assess the efficacy of a temporal sampling scheme, and 2) a simple model for identifying when a particular target is ‘observable’ due to predictable limitations, keeping in mind that there will be additional limitations from unpredictable causes which will need to be modeled as well.

2.1. The window function and its Fourier transform

Consider a sequence S of N equal time intervals, which span the entire duration T over which observations may be made. Each interval represents a sampling opportunity, during which an observation may or may not actually be made.

Consider an observable that varies sinusoidally over this total span of time, with frequency ω . If observations are made at every possible time slot, a Fourier transform of the time series observations will be a delta function. By generalization, any transient observable can be decomposed into its Fourier power spectrum, modulo the Nyquist sampling limit at the high frequency end, and the total time duration T at the low frequency end. However, if not all elements of the sampling space are filled by real observations, aliasing will occur, and side-lobes in frequency space will be seen. These can lead to specious conclusions in characterizing the variability. **The goal of optimizing observing cadences is to minimize the power in these side-lobes** in (Fourier) frequency space.

When the elements of the sequence S mentioned above are populated by unity in the event that an observation is made at the corresponding time slot, and populated by zero when an observation is *not* made, S is the so-called “window function” (WF). We designate the power spectrum of the WF (the absolute value of the Fourier transform of WF) by PWF . The observed power spectrum of a time varying observable is then its true power spectrum convolved by the power spectrum of the window function. If all elements of S are unity, for a sinusoidal variation the PWF will be a delta-function, and for more complex variation and for transient phenomena, the true power spectrum will be revealed. This is the ideal case. In practice, because we want to sample many fields, and because our the time-line is interrupted by both predictable and unpredictable limitations, S will be only partially filled, in fact, only *very sparsely* filled. **Thus the aim of time sequence planning is to make the PWF as close to a delta-function as possible** – i.e. given limits on how S can be filled, and what the filling fraction is allowed to be, what strategy is to be adopted about when to make the observations so as to minimize the power in the side-lobes of the PWF .

2.2. The implication of the *PWF*

Let us study a simple example for clarification. Consider a time-series of 1000 equal intervals. Consider the window function when observations are made in each of the 1000 intervals. The *PWF* for this case is shown in the top panel of Fig. 1, which is a delta-function. Each unit in the abscissa of this plot is $1/Nt$, where N is the number of equal intervals in the time-line (in this case 1000), and t is the duration of each time interval in real units (e.g. seconds, days, etc.). If a sinusoidal signal with period P were sampled by this sampling window, we would get a spike in the power spectrum at Nt/P . The ordinate shows the power at each frequency interval. In the second panel we show the case when observations are made at every third interval. Only 333 of the 1000 sample points have actual measurements. Note that the height of the spike at zero is reduced, and the power has been redistributed predominantly into two satellite spikes at ± 333 , each of which also has some low-level wings. With this sampling, a sinusoidal signal with period P will appear aliased, with spurious detected periodicity at P_1 and P_2 , where

$$1/P_1 = 1/P + 333/(Nt) = 1/P + 1/(3t) \quad (1)$$

$$1/P_2 = 1/P - 333/(Nt) = 1/P - 1/(3t) \quad (2)$$

Accordingly, a sinusoidal signal with a period of $2t$ will have an alias at $6t/5$ and $6t$.

The object of choosing a filling strategy is to minimize the redistribution of power into false frequencies. A commonly used strategy is to make the duration of the gaps (when there is no sampling) follow a geometrical progression. Again 333 of the 1000 sampling points are filled, but with geometrically spaced gaps. The resulting *PWF* is shown in the 3rd panel. Note how the power in the central peak is diminished even more than for the regular observing cadence, but that the redistributed power is confined to frequencies very close to the true one. An even better sampling scheme is one where the gaps are randomly distributed. Again 333 out of 1000 sampling points were filled. The results are shown in the 4th panel of Fig. 1. Note that the power in the true frequency is higher than for the geometrical sampling case and that the redistributed power does not pile up in any part of the frequency domain as aliases.

We have ascertained that given a filling fraction for sampling, a strategy where the gaps are of random duration with mean $\tau = N/n$ (where n is the number of points where observations are actually obtained) produces the least aliasing.

In practice, the filling of a time-line for a real telescope is not unconstrained. We deal next with the first level of limitations for the proposed *LSST*.

2.3. The Inherent Window Function of the projected *LSST*

Here we consider an idealized model, in which the *LSST* is to be operated for 5 years in the mode described above. The total expended time per pointing and exposure is 30s. This means that there are 2880 time-slots per 24 hour day for 1825 consecutive days, resulting in an observing time-line S that is 5.26×10^6 time-slots in duration.

We start by setting all elements of S to unity. Without loss of generality, we can start the sequence at the beginning of a night at new moon. We define a night to be 8 hours long – so the first 960 time-slots are night time, the next 1920 are day, and not suitable for observation. S is set to zero for all time-slots that are in day time. This diurnal pattern is propagated throughout the 5 year observing sequence. The moon’s phase and position are simulated with a simple model. Conservatively we estimate that observations can be made when the moon phase is within 0.2 of new moon, or when the moon is well below the horizon. When the moon’s phase and position are prohibitive, the corresponding elements of S are set to zero. We choose a target that is on the meridian at midnight on the first night, and propagate its motion at the sidereal rate from there on. Time slots where the object is more than 2 hours away from the meridian are (conservatively) deemed unobservable. It is sobering to realize from running the above model, that a given target is observable for only 5% of the total time-line!

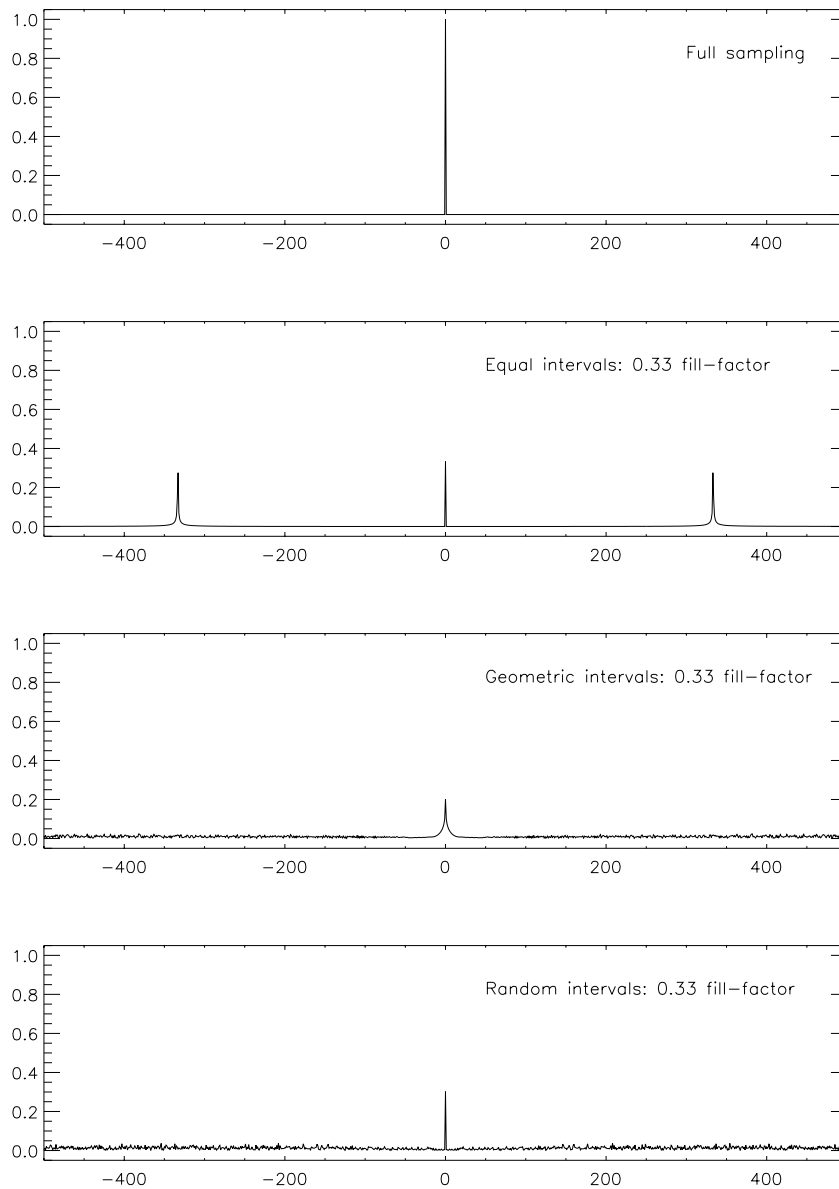


Figure 1. Figure illustrating the Fourier power spectrum of several idealized window functions. The top panel shows the power spectrum for an uninterrupted time-series. The 2nd panel for when observations are made only 33% of the time, and the gaps are regular and of equal duration. The third panel is for the same fraction of observations, but with gaps whose sizes progress as a geometric sequence, and finally in the last panel, the case for the same filling fraction but with gaps that are of random duration. See §2.2 for detailed explanation. Note that gaps in sampling that are chosen at random produce the least aliasing of the 3 partial filling cases, and most resembles the power spectrum of the uninterrupted time sequence.

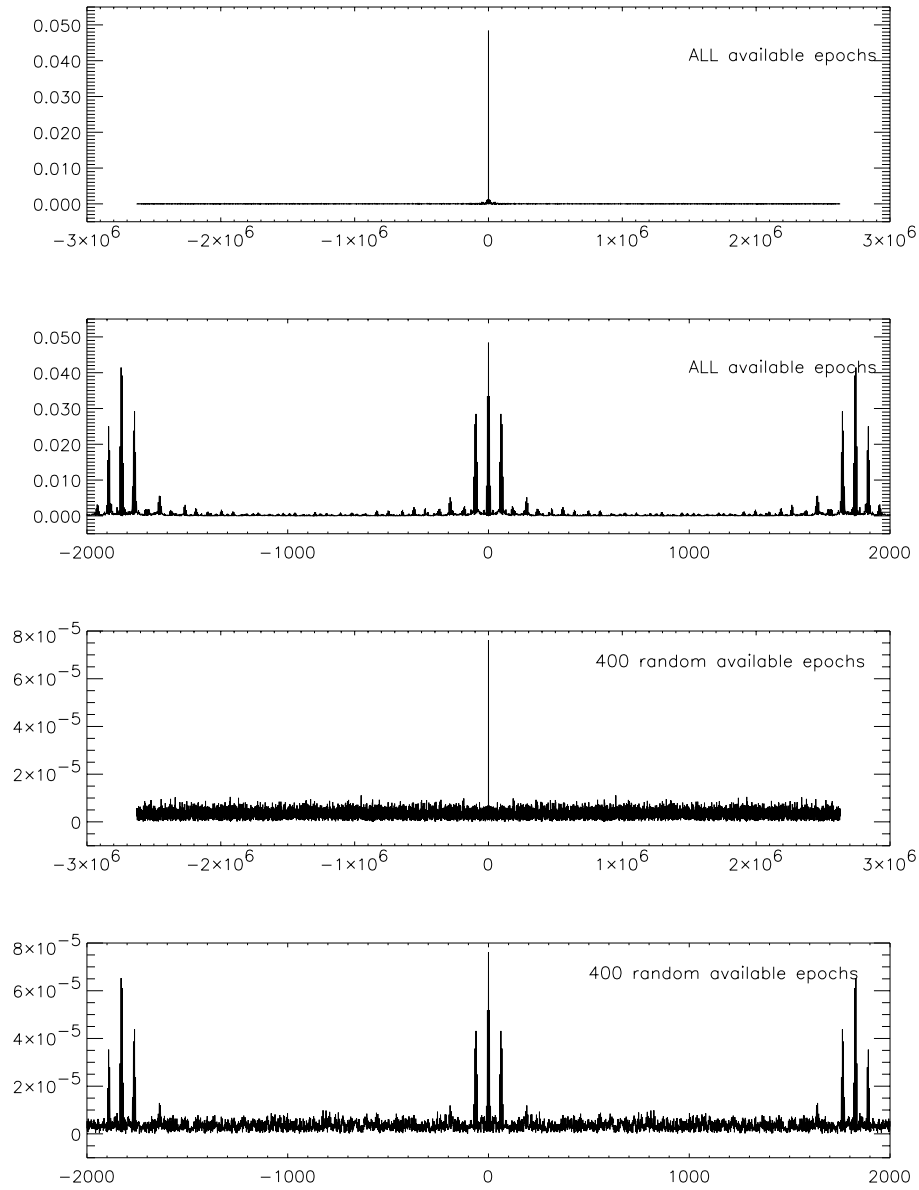


Figure 2. Like Fig. 1, except that this illustrates the power spectrum of the window function that is interrupted by diurnal, lunar and seasonal patterns that limit access to any given target. The first 2 panels (the 2nd is a blow-up of the central region of the top panel) show the power spectrum if a given target is observed at *all* epochs when it can be accessed in acceptable conditions of position in the sky and moonlight contamination. The lower 2 panels show the case when only 400 observations are made, at randomly distributed epochs when access and conditions are acceptable.

The time-line S populated according to the prescription above corresponds to the situation when a particular target is observed at *all* available time-slots. One can do no better for that target, given the inherent predictable limitations above. The window function derived from this sequence is shown in Fig. 2, and shows the intrinsic limitation. We cannot prevent this level of smearing of the timescale information and uncertainty. Since in reality we will observe a given target with much lower frequency, the real *PWF* will be worse. Our goal is to see how to minimize further damage.

The power spectrum of the inherent window function for a typical target as gleaned from the simple model described above is shown in the two higher panels of Fig 2. Note that a 5 year timeline has $N = 5.256 \times 10^6$, and $t = 30$ s, following the notation introduced in §2.2. The upper panel of Fig. 2 gives the full power spectrum, showing re-distributed power in the wings of the central peak. The second panel gives a blow-up of the same central region, showing the structure in the re-distributed power, which leads to some strong aliasing. However, since this is an inherent limitation of one telescope at one particular site, there is nothing to mitigate this situation, short of having a second telescope at a different longitude, which would reduce the diurnal alias.

Anticipating that a 5 year project will result in approximately 400 actual measurements per field, and choosing these 400 epochs *at random* from the permissible or available epochs, we can simulate what the window function for an actual data set might look like. The power spectrum for such a simulated observation set is shown in the 2 lower panels of Fig. 2, again showing the full spectrum, as well as a blow-up of the central region. Comparing with the 2 top panels, we see that the random selection of epochs has not created new artificial features in the power spectrum, even though some low level of confusion is spread across the entire spectrum.

We can surmise that by choosing observing epochs at random (about a mean rate) we do not add significant aliasing to the sampling efficacy over the intrinsic difficulties due to the diurnal, lunar, and seasonal patterns that affect access to any given target.

3. A STRAWMAN OPERATING MODEL FOR VARIABLE OBJECTS AND TRANSIENT PHENOMENA

The conclusion that choosing epochs at random from the range of possible epochs is a good strategy is a very powerful one in terms of constructing an implementable operation plan. This conclusion is made more compelling by the realization that we have thus far ignored the effects of weather and other *unpredictable* causes that can render a given epoch unusable. Were we to put in place a time-table for observing, and if the sampling efficacy were to depend on the accuracy to which we could adhere to such a time-table, the results would be vulnerable to the vagaries of weather. Clearly a dynamical procedure, where observing choices at any point in time depend on the history of what has been obtained, is needed. Such a scheme is easiest to realize if we take the freedom afforded by the random element in choosing observing epochs.

There is a difference between choosing the epochs at random before hand, and choosing them as we proceed. Most importantly, when we choose in advance, all possible epochs have equal weight, and the chosen epochs spread themselves evenly over the available space. However, if we choose the ‘next field to be observed’ by selecting at random from the available fields, we may not converge upon the same desired result. In order to have a dynamically robust operating model, we are driven to the second mode, and must ensure that for instance, all observations of field X are not piled into the last of five years! We must also keep in mind that we must group contiguous fields into zones, and schedule the zones rather than individual fields, so that *LSST* need make only small slews from one exposure to the next. As a result of these considerations, the following strawman operating model is proposed:

1. Group contiguous fields into zones of order 100 fields. Zoning along ‘constant’ Right Ascension is preferred, since it will keep schedulability criteria the same or very similar for all its constituent fields.
2. Keep a history of all observations that have already been made, both by zone and by field.
3. Construct an ‘Urgency Function’ U , that alerts us to increased urgency for re-observing a particular zone (or field). If τ is a pre-determined desired mean time between exposures, and t is the time elapsed since the last exposure, a possible form for U is $e^{t/\tau}$.

4. Choose the ‘next’ zone to be observed by drawing at random from all possible zones, but with each zone weighted by U (or replicated U times). One can also envisage weighting U by a more complicated function of the history of all prior observations, and also by a quality index indicating the quality of data. If certain spacings are to be given higher weight in deference to particular science goals, that can be accommodated as well. Note that the value of τ (above) is driven by how often we can revisit a given field. A program to search for X-ray binaries may require very frequent revisits, but may be limited to a few select fields. Thus τ may have different values for different fields.

4. CONSTRAINTS ON FILLING THE TIME-LINE

In our hypothetical 5 year operation of the *LSST*, how many epochs will we have if we are to cover the whole sky? From our strawman ‘observability’ model described in §2, we deduce that we can actually observe for about 14,000 half minute time slots, given the diurnal and lunar constraints. However, a *given* target cannot be observed all of the time: a target that is within an hour of the meridian at midnight during a particular lunation can be observed during any of 9,300 slots. If it has an hour angle up to 3 hours at midnight, the number of opportunities decreases to about 6400 slots, if up to 5 hours over at midnight, it has 3,400 slots, and if it is as much as 7 hours over at midnight it has only 235 slots. At the equator, a 2 hour section of sky covers about 3500 square degrees, or 500 *LSST* fields, and serves as an order of magnitude estimate.

This tells us that any 2 hour sector is observable for 5 consecutive lunations. Each month one such sector is retired, and a new one is available. At any one time 5 such sectors, or up to 2500 fields are ‘active’. If all active fields are observed equally, in any lunation a given field can be observed at most 5 or 6 times. Additionally, there will be losses to weather and other inclemencies. If we want a larger number of epochs per month, we must either:

1. Decide not to cover the whole sky, or
2. Do some sectors more often than others, or
3. Selectively tune the frequency of different fields/zones/sectors according to scientific goals.

What does this imply about the proposed 9 visits to each target per lunation preferred by the NEA project? First a 2 hour section of the sky, which comprises 500 *LSST* pointings, cannot be fully imaged even once during the lunation when the section is 3 months from optimal (where optimal is when the target is near the meridian at midnight). It can be imaged as many as 7 times if it is 2 months from optimal (which falls short of the current model preferred by NEA hunters). Within a month of optimal transit, all 9 pointings can be accommodated. Other more economical schemes may do considerably better, but the message to take away is that the trade-off space is tight, and it is imperative to optimize the NEA search strategy by balancing search goals against amount of sky coverage and number of exposures per field per lunation.

There is another ramification of the proposed NEA cadences. Consider, for example, the case of the 3 observations in the same night spaced by a few hours, and 2 more such sets of 3 in the same lunation. Such patterned observing cadences will produce aliasing in the power spectrum of the window function, possibly rendering difficult the search for time varying phenomena with certain ranges characteristic time-scales. This can be ameliorated, for example, by running the *LSST* project for a longer duration, and in the outlying years, with the NEA project completed, we fill in the sampling deficiencies by appropriately changing the timing strategy. On the other hand, the other proposed strategy for NEAs mentioned in §1 requires a less regular pattern (days 1, 2, and 10), which yields lesser aliasing hazards.

5. SUMMARY AND CONCLUSIONS

This article brings forward some of the global considerations in designing an operating plan for the *LSST*. We have presented a study of the temporal sampling that *LSST* can provide, and have compared some of the timing sequences and schemes that minimize aliasing. More cases were actually studied than can be presented here.

An important result is that selecting observing epochs at *random* produces the smallest aliasing effects. This is particularly useful, because random sampling is more operationally robust than a time-table driven observing scheme.

The diurnal, lunar, and seasonal patterns that govern when a target field is actually observable produces distinctive aliasing. However one cannot improve upon this without building a second telescope at a significantly different longitude. Science programs should be tested to see if the aliases produced by these natural rhythms are fatal – mostly these will affect time-scales of order a day.

The total time available, if distributed equitably among all targets, is only sufficient to observe each target about 5 times per lunation. This is *before* we account for weather and other losses. Since many projects, most notably the NEA search and SNe search programs have indicated need for more frequent sampling and/or multi-passband observations, it is not going to be possible to survey the whole sky, nor perhaps to accommodate all investigations. Using the *LSST* for followup observations will also be at a large cost to other goals.

Observing cadences required or demanded by one project can have a deleterious impact on other projects. Major trade-offs, either in scientific priority, or in the fraction of sky covered will be necessary.

In light of this we are driven to make a few, now obvious, recommendations:

- The science goals of the *LSST* need to be prioritized as early as possible.
- The observing cadence and sky coverage fraction for the NEA search can impact other science projects heavily. An understanding of various NEA search strategies and their respective relative impact on the NEA search program is a key ingredient towards arriving at an overall *LSST* operating plan. Necessary simulations to arrive at this understanding are needed as early as possible.
- Consider an auxiliary followup facility, preferably involving more than one telescope and situated at sites with different longitude.
- Consider a longer project life for *LSST*, and block science projects with similar sampling requirements into different years of operation.
- There is need for a precursor experiment to validate the predictions made here and to explore alternatives.

It is our hope that the tools of evaluation and line of reasoning presented here will be useful in deliberating the science priorities and eventual operating plan for the *LSST*.