Towards a Design Reference Mission for the Large Synoptic Survey Telescope

The Science Working Group of the LSST, Michael A. Strauss, chair

ABSTRACT

Three nationally endorsed decadal surveys have stated that a high priority for US planetary science, astronomy, and physics over the next decade should be a dedicated wide-field imaging telescope with an effective aperture of 6–8 meters. Such an instrument, in dedicated survey mode, could catalog 90% of the asteroids whose orbits cross that of Earth with size greater than 200-300 meters, revolutionize our understanding of the Kuiper Belt with orbits of over 10^5 objects beyond the orbit of Neptune, allow unprecedented mapping both of the distant Galactic halo and the local solar neighborhood, explore the optically faint variable universe, obtain huge samples of supernovae to redshifts of unity, and yield insights into the dark energy by measuring the gravitational lens signal with high accuracy. This document expands upon this scientific case, and concludes that building such a facility is indeed in the best interest of the US scientific community. We briefly discuss two different approaches that have been suggested to implement this goal: a single monolithic telescope with an 8.4m primary and a 3-3.5° field of view, and an array of 1.8m telescopes, each with a large field of view and a separate imaging camera. Deciding the superiority of either approach is left for future work.

1. Introduction

Many of the dramatic advances in astronomy have come about via massive surveys of the sky, in wavebands from gamma-rays to radio. The ability to carry out these surveys, in turn, have been driven by technological advances: in telescope design, detector technology, and software and processing power. For decades, the state of the art in wide-field imaging in the visible part of the spectrum was the photographic Palomar Observatory Sky Survey (POSS), and its southern counterpart carried out in Australia. However, with the advent of large-format CCDs, and the electronics to build cameras with substantial numbers of these devices, it has been possible to build wide-field imaging cameras using these much more sensitive detectors. As of this writing, many of the 4–8 meter class telescopes around the world have CCD mosaic imaging cameras with fields of view from 20 arcminutes to over a degree, and many more are planned; moreover, these instruments are increasingly being used for massive surveys of the sky (such as the NOAO Deep Wide Field Survey[1], which covers 18 deg^2 in BRIJHK, the Deep Lens Survey[2], which will cover 24 deg^2 in four bands, and the planned Legacy survey on the CFHT[3] which will cover 1300 square degrees in three bands). In addition, smaller telescopes are being used to carry out dedicated variability surveys, such as the OGLE[4], which uses a 1.3m telescope and an imaging camera 35′ on a side to obtain lightcurves for literally millions of stars in the Galactic Bulge and the Magellanic Clouds, and the LINEAR[5] survey, which uses a pair of 1m telescopes with two square degree fields of view to search for asteroids, especially those whose orbits may take them close to the Earth.

---

[1] `http://www.noao.edu/noao/noaodeep/`

[2] `http://dls.physics.ucdavis.edu`

[3] `http://www.cfht.hawaii.edu/Science/CFHLS/left.html`

[4] "Optical Gravitational Lens Experiment"; `http://bulge.princeton.edu/~ogle`

[5] "Lincoln Near Earth Asteroid Research", `http://www.ll.mit.edu/LINEAR`

The current state of the art survey for wide-field optical imaging is the Sloan Digital Sky Survey (SDSS; `http://www.sdss.org`; York et al. 2000) which uses a dedicated 2.5m telescope with a 3° diameter field of view to obtain photometry to $r_{AB} \sim 22.5$ in five photometric bands ($u\,g\,r\,i\,z$). As of this writing, it has imaged roughly 7000 unique square degrees of mostly high-latitude sky in four years of operations. Its principal scientific drivers are related to the large-scale distribution of galaxies, but the scientific results that have come from these data range from studies of the colors of main-belt asteroids (Ivezić et al. 2003) to galaxy-galaxy gravitational lensing (McKay et al. 2001) to the discovery of the highest-redshift quasars (Fan et al. 2003). The lesson to take away from this is that wide-field survey data, even if focused on a specific scientific goal, will find wide applicability to a very large range of astronomical topics.

A useful measure of the surveying capability of an imager on a given telescope is its *étendue*, the product of the solid angle Ω subtended by the camera, and the collecting area A of the telescope. The SDSS has an étendue of $A\Omega = 5.6\,\mathrm{m}^2\,\mathrm{deg}^2$; for comparison, the wide-field Suprime-Cam on the 8-m Subaru telescope has $A\Omega = 13.5\,\mathrm{m}^2\,\mathrm{deg}^2$. Such étendues, impressive though they are, are simply not adequate to address some of the most pressing scientific questions that face us. Among these, in order of characteristic distance from the Earth, are:

- The population of Near-Earth Asteroids. Ever since the realization in the late 1970's that the great extinction of the dinosaurs at the Cretaceous-Tertiary boundary 65 million years ago was probably caused by an asteroid or cometary impact, the dangers of future impacts has been studied intensely. There does exist a population of asteroids whose orbits take them close to Earth; a collision with one as small as one kilometer could be devastating to life on Earth, and even one as small as 200-300 meters would cause wide-spread death and destruction. The US Congress has mandated that at least 90% of the Near-Earth Asteroids with diameters greater than one kilometer be discovered and orbits determined for them by the year 2008. It is now uncertain whether this goal will be met (Jedicke et al. 2003), even with the best efforts of surveys like LINEAR. The enhanced goal of discovering and tracking most of the asteroids above 200 meters will clearly take a telescope with much larger collecting area.

- The study of Kuiper Belt Objects. These are asteroids with orbits beyond that of Neptune. The first of these was discovered only a decade ago (Jewitt & Luu 1992); roughly 800 examples of this major new component of the solar system are known now. Much larger samples, which will require massive deep surveys of large areas of sky, would have much to teach us about the dynamical history of the solar system, the origins of comets, and formation mechanisms of planets.

- The study of the variable sky. As Paczynski (2000) has remarked, our samples of variable objects are massively incomplete at all flux levels, on a variety of timescales. This is especially true at faint flux levels (say, below 20th magnitude), where very little is known about variability. Microlensing surveys such as OGLE and MACHO have shown the richness of information available in densely sampled and accurately calibrated wide-field photometric data, and the range of scientific problems such data can address.

- The structure of our Milky Way. The detailed distribution in position and velocity space of stars of different ages and metallicity in our Galaxy encodes information on how the Milky Way has formed and evolved. One of the most important discoveries along these lines over the last five years, made possible with the current generation of wide-field imagers, is the detection of substructure in the Galactic halo, due to the tidal stripping and disruption of globular clusters and satellite galaxies as they orbit the Milky Way. A truly global view of this process will require a deep wide-field multi-color imaging survey; synoptic observations will allow proper motions to be measured as well. Such data will also allow parallax to be measured for an

essentially complete sample of stars within 100 pc or so, allowing a definitive determination of the stellar initial mass function and temperature-luminosity relation at the faint end.

- The distribution of dark matter in the universe, and its evolution with cosmic time. Weak gravitational lensing by large-scale structure has long been recognized as a tool for probing the distribution of dark matter. Weak gravitational lensing in the field has now been detected by a variety of groups, albeit still at only moderate statistical significance. The lensing distortion of a given galaxy is due to a weighted integral of all density inhomogeneities along the line of sight. With photometric redshifts of these galaxies, one can hope to probe the statistics of dark matter clustering as a function of redshift; the clustering signal evolves with redshift in a way that depends on the underlying cosmological model. In particular, the measurement of weak lensing has the potential to constrain models for dark energy, such as the w parameter of quintessence and its possible evolution.

- Studies of Type Ia supernovae have led to perhaps the most exciting cosmological discovery of the last decade, namely that the expansion of the universe is accelerating. Exploring this result in detail requires large samples of supernovae to test for systematic effects, to study the physics of the supernovae themselves, to look for variations in the acceleration rate with direction in the sky, and so on.

These and many other forefront scientific issues can be addressed with a dedicated telescope facility of étendue of order 250 meters^2 degrees^2, an order of magnitude larger than what is currently available. Moreover, data taken for one of these projects will be of value for others: the multiple exposures needed to go deep on the weak lensing study, for example, can be used to search for variable and moving objects such as asteroids.

With these considerations in mind, three different national committees (all convened by the National Academy of Science) have listed as one of their top scientific priorities that a national facility be built with these characteristics. These committees and their reports include:

- **Astronomy and Astrophysics in the New Millennium**, a report written by the McKee-Taylor Committee, which assessed the need for new facilities and initiatives in astronomy and astrophysics in the decade of 2000-2010. The committee, convened by the National Academy of Sciences and involving broad input from literally hundreds of astronomers, emphasized most of the science drivers listed above. Its report was published by the NAS in 2001, and can be found at `http://www.nap.edu/books/0309070317/html`.

- **Connecting Quarks with the Cosmos: Eleven Science Questions for the New Century** was a report of a committee, chaired by Mike Turner, to examine broad scientific needs on the interface between astronomy and physics, to understand the physics of the universe as a whole. It put particular emphasis on the measurement of weak lensing as a probe of the dark matter distribution and its evolution. This report, published by the NAS in 2003, can be found at `http://www.nap.edu/books/0309074061/html`.

- **New Frontiers in the Solar System: An Integrated Exploration Strategy** was the report of a committee, chaired by Mike Belton, of the priorities for new facilities and missions for the next decade in solar system science. The LSST concept was endorsed specifically for its ability to discover near-earth asteroids and Kuiper Belt Objects, as well as to study the population of long-period comets. This report, published by the NAS in 2003, can be found at `http://www.nap.edu/books/0309084954/html`.

To further develop the case for such a national facility, the National Optical Astronomical Observatory (NOAO) convened a Science Working Group (SWG) to refine the scientific drivers for a Large Synoptic Survey Telescope (LSST), and to assess whether these drivers will remain compelling over the timescale that such a facility

could be built. The current document is the first report of this committee, focused mostly on the former question. The SWG membership (Appendix D) was picked from applications from the US astronomical community, and includes experts in all areas of science outlined above. Note that because the NOAO itself has strong interest in being involved in the building of any implementation of LSST, no members of the NOAO staff were included in the SWG, to avoid any possible conflict of interest. However, a number of members of the NOAO scientific staff did work closely with the SWG in their discussions.

The SWG has met a variety of times both face-to-face and by phone since it was convened in September 2002. All its meetings have been open to the scientific public. It has also carried out much of its business via a publicly archived e-mail exploder[6]; as of this writing, 53 people (the majority of course not members of the SWG) subscribe to this mailing list.

In this report, we use the term, Large Synoptic Survey Telescope, or LSST, to refer to the generic notion of a telescope or array of telescopes with a large étendue (i.e., of order 200-400 deg^2 meters^2), to carry out the scientific tasks outlined above, and described in more detail below. As we describe in appendices to the report, there are currently two specific implementations of the LSST concept which have been proposed and are undergoing design studies. The first is a single monolithic telescope, with a primary of diameter 8.4 meters (but effective aperture of 6.9 meters due to its large central hole) and a three-mirror design to allow a large field of view. This concept has been referred to in the past as the Dark Matter Telescope (or DMT, in reference to its ability to measure the distribution of dark matter through weak gravitational lensing), but is more generally known now as the LSST (see `http://www.lsst.org`, and Appendix B). In the present report, we reserve the term "LSST" without adornment to refer to the generic idea of a large étendue system, and use the term "8.4m LSST" for the specific monolithic 8.4m design.

Another approach is that of the Panoramic Survey Telescope and Rapid Response System (Pan-STARRS, `http://pan-starrs.ifa.hawaii.edu` and Appendix C) which uses an array of wide-field, moderate size telescopes (1.8 meters in diameter). The Pan-STARRS team has funding in hand to build a system of four such telescopes and hopes to see first light for the first of these telescopes in 2006. 15-20 such telescopes with wide-field detectors would reach the étendue called for in the decadal surveys mentioned above; we refer to such a system as "Multi-STARRS".

The present document focuses on explaining the science case for the LSST, and does not discuss the relative merits of a single aperture design like the 8.4m LSST and a distributed aperture like the 20-telescope incarnation of Pan-STARRS; this is future work for the LSST SWG, as discussed in further detail in § 12. A summary of the basic parameters of 8.4m LSST, Pan-STARRS, and other related programs is given there as well.

[6] `http://www.astro.princeton.edu/~dss/LSST/lsst-general/INDEX.html`

2. The Science Justification

The science cases detailed below describe the abilities of the LSST to make substantial advancements in various fields. Throughout, we have assumed that the LSST would have an étendue of 250 m^2 deg^2 and the ability to obtain well-sampled images to $r \sim 24$ $(10\,\sigma)$ in good seeing in a single exposure of 10-30 seconds, and would be a dedicated facility operating for of order a decade.

Given the long timescale required to build any LSST implementation, it is difficult to predict the state of these fields at the time LSST will see first light. Nevertheless, we attempt to determine, based on other current and planned facilities, the extent to which the science will remain compelling. More generally, it is possible, indeed likely, that the most exciting science that LSST may be capable of is in areas which we did not even consider. Ten years ago, many of the current frontiers in astrophysics (the discovery of dark energy, the ubiquity of supermassive black holes in ordinary galaxies, and the discovery of planets around other stars, to name a few) were not anticipated. It will surely be amusing when, in 15-20 years time, we compare the scientific output of LSST with the naivety of this report. However, by focusing on specific scientific goals and designing a telescope facility and survey to meet those goals, one generates a dataset capable of addressing a wide range of (often unanticipated) scientific issues, as survey after survey have shown.

We cover the scientific areas in the order in which they were presented in the introduction: Near Earth Asteroids, Kuiper Belt Objects, the variable universe, Galactic structure, weak lensing, and supernovae, followed with a brief (and very incomplete!) discussion of various other areas of science which LSST can address. Three concluding chapters describe how these various programs may be integrated, some of the issues that our deliberations raised, and a discussion of data distribution and data rights issues. Appendices include descriptions of a possible observing cadence, descriptions from proponents of each project of the 8.4m LSST and Pan-STARRS, and the membership of the SWG.

3. Potentially Hazardous Asteroids

3.1. The Impact Hazard

The Earth is immersed in a swarm of Near Earth Asteroids (NEAs) whose orbits approach that of the Earth. A subset of about 20% these, so-called Potentially Hazardous Asteroids (PHAs), are in orbits that pass close enough to the Earth's orbital path (< 0.05 AU) that planetary perturbations on a time scale of a century or so can lead to intersecting orbits, hence a non-zero possibility of collision. It seems only prudent to discover and monitor such asteroids in order to be sure that no impact is in our immediate future – or to identify any object that is in fact going to collide in the next century or so.

The relative orbital stability of PHAs, even those that are Earth-crossing, makes discovery and cataloging a practical task. NEAs are identified from their motions against the background of stars using automated wide-field astronomical telescopes. Today such discovery telescopes are typically of about 1 m aperture and can detect NEAs as faint as magnitude $V = 20$. Follow-up observations are generally made using other telescopes, often by dedicated amateur observers. Astrometric observations stretching over a few weeks of time are usually sufficient to define an orbit and permit close approaches to be predicted out to a century or more in the future. The current Spaceguard Survey has as its goal the discovery by 2008 of 90% of the NEAs larger than 1 km diameter, together with determination of orbits sufficient to ensure no collisions with the Earth for at least the next century. If a potential impactor is found, there is likely to be sufficient time for society to develop mitigation plans (e.g., Schweickart et al. 2003).

In a 1992 NASA report, "The Spaceguard Survey Report" (D. Morrison, ed.), primary emphasis was placed on the potential of large impacts to cause global damage. Studies (e.g. Toon et al. 1997) have indicated that the threshold for global environmental damage occurs for an impact energy of roughly 1 million megatons, corresponding to an asteroid approximately 2 kilometers in diameter. An impacting asteroid larger than about 2 km in diameter (but possibly as small as 1 km) would lead to a major climatic disaster, similar in nature to "nuclear winter", producing a year or more of disruption of agriculture that could lead to the death of a significant fraction of the world's population, mostly from starvation.

There are about 1,000 NEAs, hence about 200 PHAs, of diameter > 1 km. The frequency of impact of such bodies is estimated to be about once in 500,000 years. A simple division of casualties of order a billion divided by a time between events of half a million years suggests a "fatality rate" from such events of more than a thousand per year, perhaps several thousand. The time-averaged risk is greater for smaller, more frequent events, down to the lower limit of size for global effects (cf., Figure 2 discussed below), thus the major source of uncertainty in quantifying the risk is the minimum size at which a global catastrophe would result.

The damage from impacts by sub-kilometer objects is confined to the area of impact, but the explosions from such asteroids striking at cosmic velocities are still of a magnitude that dwarfs other more familiar natural catastrophes. The Tunguska impact of 1908 in Siberia was due to an asteroid of only about 60 m diameter which penetrated to within a few kilometers of the surface before explosively disintegrating in the atmosphere. The resulting airburst is estimated at 10-15 megatons energy, and it flattened a region of forest of more than 1000 square kilometers, the size of a major city. Stony asteroids larger than about 150 m diameter are able to penetrate to the surface and form craters. The 3/4 mile-wide Meteor Crater in Arizona was caused by a 15-megaton ground blast from a relatively rare iron projectile, which can reach the ground at much smaller size. The smallest known impact crater on the Earth caused by a stony impactor is about 2 km in diameter.

Since most impactors will strike the ocean, it is important to consider the hazard to coastal communities from impact tsunamis. The process of wave formation and motion across an ocean allows the impact energy to propagate much farther from the point of impact than does the blast wave from an airburst or land impact, so

that shorelines thousands of kilometers distant may be at risk. The process of formation of impact tsunamis tends to produce much shorter wavelengths than those of seismic tsunamis, with associated reduction in the expected run-in and run-up of the wave as it reaches shore. As discussed further below, the hazard from tsunamis appears to be less than had been feared in some earlier studies, but it still exceeds that from ground impact by sub-kilometer asteroids.

This discussion of the impact hazard deals with impacts of asteroids: projectiles composed primarily of rocky and metallic materials. The icy, long-period comets, however, constitute another potential population of impactors; one often refers to the population of near-earth asteroids and comets jointly as "Near-Earth Objects" (NEOs). Recent studies indicate that the number of comets relative to asteroids drops off sharply toward smaller sizes, and that comets in the sub-kilometer range are essentially absent from the inner solar system; thus the impact hazard from comets is quite a bit less than that for asteroids. This is a good thing, as comets tend to be on highly elongated orbits with periods of hundreds to thousands of years, meaning that early detection of hazardous long-period comets is beyond current technical capabilities of optical telescopes. Thus in this report we consider only the impact risk from asteroids, focusing on NEAs and not the more inclusive class of NEOs.

The challenge to a NEA survey of the type discussed here is to identify a large number of individual objects that could pose a threat of impact. While an extensive survey will yield a great deal of scientific information on the NEA population (as well as on other populations of solar system objects), it is not our goal to improve impact statistics. We already understand the impact population well enough to provide order-of-magnitude estimates of the hazard and to compare this risk with other natural hazards, and remaining uncertainties are dominated by our lack of understanding of the aftermath of an impact, not by population statistics (see discussion below). The object here, from the perspective of protection of our planet, is to find the individual objects, one at a time, and calculate orbits of sufficient precision to determine whether each one is or is not a threat. The question of primary concern to society is not the probability of impact, but whether (and if so, where and when) such an impact will occur within our grandchildren's lifetime. One objective of the LSST is to answer this question.

The strategy currently pursued by the Spaceguard Survey of cataloging NEAs larger than 1 km will be continued in an expanded survey carried out with the LSST instrument, but with vastly greater discovery rates reflecting the rapid increase in number of asteroids at smaller sizes. One significant difference is that the LSST will generate its own follow-up astrometric data, since the objects being discovered are both too faint (magnitude $V = 24$) and too numerous to rely on volunteer follow-up with small telescopes. The product of such a survey is not just the discovery of asteroids but the determination of their orbits with sufficient precision to ensure that they will not threaten the Earth within the next century or so. The vast majority of the asteroids discovered can be rejected as threats very quickly, while a handful will require further study to refine orbits. Some fraction will also prove to be suitable for physical studies by other techniques including radar, which is a powerful technique for both imaging and precision determination of orbits. If any NEA is found with a significant probability of impacting Earth, the warning is likely to be several decades. Given such advance notice, it reasonable to expect society either to find ways to deflect the orbit so the asteroid will not hit, or otherwise to mitigate the danger, for example by evacuating the target area for a small asteroid predicted to hit in an area of low population.

3.2. The Hazard of Subkilometer Impacts

NASA has recently commissioned a study to evaluate the current state of knowledge of the population of NEAs (PHAs), the frequency of impacts, the nature of damage caused as a function of size, and what can or should be done to address the hazard (Near-Earth Object Science Definition Team 2003; hereafter the *NASA report*). Figure 1 is a plot, taken from that report, of the currently estimated population of NEAs as a function of

size.

In the NASA report it is concluded that there are about a million NEAs larger than $\sim$ 50 m in diameter, the smallest size capable of causing significant ground damage. The frequency of such events, similar to the Tunguska event over Siberia in 1908, is once in several hundred to a thousand years. At the larger end of the size spectrum, an impactor larger than about 1-2 km in diameter is likely to cause a global climatic disaster, which could lead to billions of deaths, mainly through starvation due to disruption of agriculture for a year or more. The frequency of such events is of the order of once in 500,000 years.

This recent NASA report focuses on quantifying the risk from subkilometer NEAs, since it is now technically feasible to consider cataloging most of the smaller ones. To make a sound cost-benefit decision requires a more careful evaluation of the impact hazard from smaller bodies. Figure 2 presents a summary of the recent evaluation of impact risk.

At the lower end of the size spectrum, asteroids smaller than about 50 m in diameter are effectively stopped by the Earth's atmosphere, so no ground damage occurs. In the size range of 50–200 m diameter, impactors explode low in the atmosphere with energy comparable to a large nuclear explosion (tens to hundreds of megatons), causing extensive ground damage over a radius of tens to hundreds of km. Over the ocean, such an event would be harmless. As can be seen in Fig. 2, damage from land impacts or airbursts is very much less than from larger impactors ($>$ 1 km) that can cause global damage.

Impactors larger than about 150 m in diameter hitting the ocean may generate potentially devastating tsunamis. Based mainly on work by Ward & Asphaug (2000), the recent NASA report gives estimates of annualized risk of about 180 affected people/year, peaking in a size range of 200 to 500 m diameter of impactor. In the case of tsunamis, it is expected that most affected individuals will be able to escape by evacuation, so the number of affected individuals is more a proxy for property damage than a fatality estimate. Nevertheless, this is a substantial number corresponding to property loss of tens to hundreds of millions of dollars per year. As with land impacts, the magnitude of destruction from tsunamis increases more slowly with increasing event size than the frequency decreases, thus the risk spectrum peaks between 200m and 500m (Figure 2). There are estimated to be $\sim 50,000$ NEAs ($\sim 10,000$ PHAs) larger than 200 m diameter, with an impact frequency of about once in ten thousand years (Figure 1). Thus impact tsunamis are expected no more frequently than once in ten thousand years, but could affect a million individuals if one were to happen in the present era.

However, there remains a large uncertainty in the actual hazard from tsunamis. Figures 2 and 3 show that the range in estimated hazard from impactors 150 to 800 meters diameter (dominated by the tsunami risk) is nearly an order of magnitude, from a minimum of about 60 affected individuals per year to a maximum of 380 per year, with a "nominal" value of 180 per year. Most of this uncertainty is due to our poor understanding of the propagation of impact-generated tsunami waves, particularly in the way they propagate onto continental shelves and onto land. There is further uncertainty as to the human population that is vulnerable to impact tsunamis, as most heavily populated shorelines are protected in harbors or estuaries and not exposed to open seas. The vulnerability of this fraction of the shoreline population is particularly dependent on how far impact tsunami waves penetrate into such protected areas.

Impact tsunami waves are expected to have a wavelength comparable to the ocean depth, that is, a few to ten km. This is much longer than wind-driven ocean swells, but much shorter than typical seismic tsunamis, so there are essentially no observational data available and one must rely on theoretical results from hydrodynamical simulations. Unfortunately, this wavelength lies in the transition range between "deep" and "shallow" wave propagation, so scaling from one regime or the other is questionable. It is clear that further research in tsunami propagation is needed.

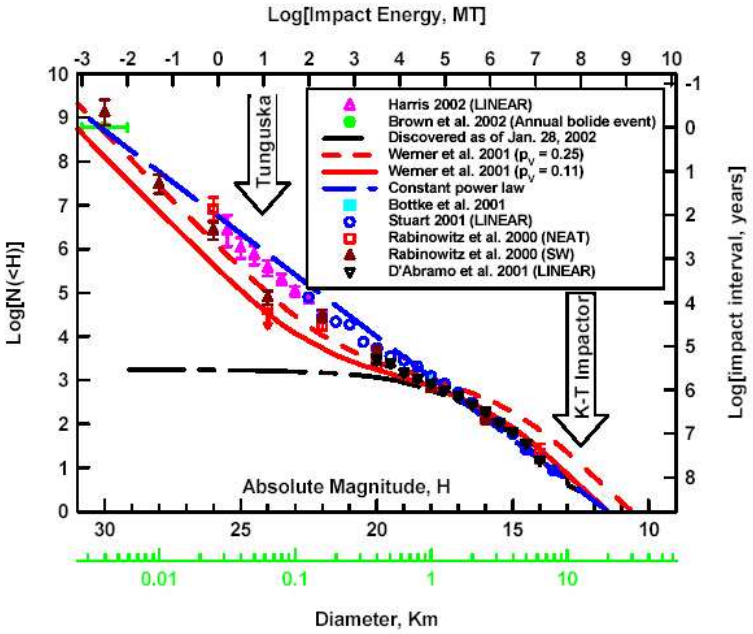

Fig. 1.— Recent estimates of the cumulative population of NEAs over the size range of potential hazard. The dashed straight line is a simple power law approximation that matches most population estimates within a power of two over the entire size range. The principal scales are cumulative number $N(< H)$ and absolute magnitude H. Also given on the right vertical scale is impact frequency, and horizontally, across the top, impact energy, and across the far bottom, diameter, based on an assumed mean albedo. This is Figure 2-3 in the NASA report.

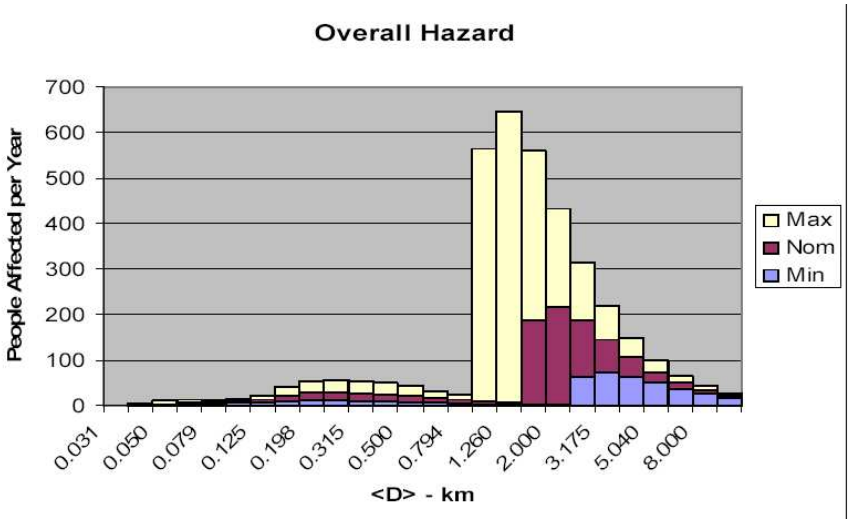

Fig. 2.— Estimates of the impact hazard from the entire NEA population before any survey is done (pre-Spaceguard). Three estimates are given, spanning the range of expected uncertainty. The total risk over a given size range is the sum of the numbers in each individual histogram over the size range of interest. As concluded in the Spaceguard Survey report, the greatest risk is from the largest impactors, which would cause a global climatic disaster. The midrange bump around 300 m diameter is from tsunami risk, and the tail extending to 50 m diameter is from impacts or airbursts over land. This is Figure 3-10 from the NASA report.

3.3. Status of Current Surveys and Next Generation Goal

Beginning in 1998, NASA established the goal to discover 90% of all NEAs of diameter > 1 km within ten years. Current surveys employ several telescopes in the size range of about 1-meter diameter, and reach to a limiting V magnitude of about 20. Now halfway through that decade, more than half of the large NEAs have been discovered. However, the discovery rate is asymptotic as completion is approached. If discoveries were a random process, then completion vs. time would be an exponential function approaching 100%, and the time to reach 90% completeness would be 3.3 times longer than the time to reach 50%. In fact, the situation is substantially worse than this because some NEAs (the ones yet to be discovered) are more difficult than others (the ones already discovered) because of longer than average orbit periods, near-commensurability with the Earth's orbit period, etc. For these reasons, the present rate of discovery is not expected to reach the 90% completion mark by 2008.

Figure 3 is a plot similar to the previous one, but showing the residual risk expected at the end of the decade of the Spaceguard Survey (2008). This estimate from the NASA report is quite conservative, assuming no increase in survey capability for the second half of the decade-long survey. They estimate that the integral completion for objects larger than 1 km in diameter, already about 60%, will be only 70% by 2008. In fact, new telescopes have come on-line during the first half of the survey and additional capability and improvements in survey technology are continuing to improve the quality of the present surveys, so it is likely that the residual risk will be substantially less than shown in Figure 3. However, the reduction will be concentrated at the large end of the size distribution, so perhaps the greatest uncertainty in what will remain for the next generation survey is whether there will still be a significant number of undiscovered objects capable of causing a global climatic catastrophe. As can be seen in Figure 3, there is at least a two order-of-magnitude range from the "minimum" to the "maximum" estimate of residual risk from large objects, mainly due to uncertainty in the threshold for global effects. Considering that the survey completion expected by 2008 is likely underestimated, and that it will likely be later than 2008 before the LSST comes on line, it appears that the residual risk from large NEAs will at that time be in the same range as that from small NEAs, that is in the range of 100-200 lives per year, worldwide.

As the threat from large impactors is reduced to about the level from smaller impacts, it is natural to consider whether it is worthwhile to continue the survey to smaller sizes, and if so what should be the next goal, and what is needed to achieve it. The NASA report concludes that a reasonable next-level goal should be to reduce the residual hazard by an additional order of magnitude, which would require discovering about 90% of PHAs down to about 140 m in diameter. Because the risk spectrum from sub-km diameter impacts is peaked toward the lower end of the size spectrum, a categorically more ambitious survey is called for, not just continuing the present surveys for another decade or two, although the latter would substantially reduce the remaining risk from large impacts.

Fully meeting the next generation goal proposed in the NASA report is probably beyond the capability of LSST or any other currently contemplated survey. However, a survey covering about 9,000 square degrees of sky along the ecliptic, three or four times a month, to a limiting V magnitude of 24.0, should come fairly close to meeting the goal. It is estimated that this search area and detection threshold should achieve a ten-year completion of about 90% of asteroids larger than about 250 m diameter, and about 80% completion down to the 140 m diameter recommended by the NASA report. Thus the LSST survey could reduce residual risk of NEA impacts by about 80%, which is the objective we adopt for the present document.

3.4. Requirements for an effective PHA survey

Preliminary evaluations of survey strategy indicate that the most effective survey pattern is to cover the ecliptic region of the sky to as small elongation as practical, preferably down to within 60° of the solar direction.

This amounts to an ecliptic band 240° long. It appears that a band as narrow as 20° wide (±10° in latitude) might suffice, although a wider band, perhaps 45° wide, would certainly do very well. Appendix A describes a "Possible Universal Cadence" for LSST observations that would satisfy the PHA survey requirements. This cadence allows a region of sky about 40 degrees in latitude by 80 degrees in longitude to be covered in a single 9 hour night. In three nights, the longitude coverage would extend to 240 degrees, thus covering all but the ±60° of longitude closest to the sun. In practice, some of the most extreme latitude/longitude regions probably could not be covered because of airmass restrictions. It also may prove to be more efficient to cover somewhat less range of latitude but cover the search region three times in the night, which could push the detection threshold a bit fainter, due to intrinsic lightcurve variations of targets near the threshold and other variables that could result in only two of three images coming up over the threshold S/N. These are matters to be studied in more detail in the future, but for now the general search strategy is clear.

The typical rate of motion of NEAs detected will be hardly greater than regular main-belt asteroids (MBAs), about 0.5 degrees/day, as they are typically seen at distances approaching an AU. Even at 1.0 degree/day, 0.7 arcseconds of trailing takes 17 seconds of time, so two successive ten-second exposures will have minimal trailing loss. One could minimize trailing loss by shifting the images before adding by the mean expected rate of motion in the area of sky being imaged. This could reduce the mean residual motion enough to allow up to 30 seconds to obtain the two images as described above without trailing loss for most targets. However, pairs of exposures must be kept to about 30 seconds total duration. This time limit sets a restriction on usable filter bands to reach an equivalent threshold of V = 24.0 for the PHA survey. Table 1 lists the integration times required to reach a S/N = 5 detection of the magnitude equivalent to $V = 24.0$ in other color bands for objects of the average color of asteroids, i.e. $B - V = 0.8; V - R = 0.4$, and $R - I = 0.4$. The last row is for a broad-band filter spanning both V and R. The first column of exposure times scaled to the effective aperture of the 8.4m LSST from the current performance of the CTIO 4-m telescope and Mosaic II imager (both calculations assuming 0.7" seeing), and the second column is for the target performance of the 8.4m LSST[7]. The difference is mainly due to assumed lower loss optical surfaces and higher QE for the 8.4m LSST. It can be seen that for the optimistic performance of the LSST, any of the four color bands can reach the target limit of V = 24.0 in about ten seconds or less. For the more pessimistic KPNO scaled numbers, any but the B band can do it. Extremely short times (e.g. V+R) are not too helpful for reducing time to cover area since that rate becomes dominated by readout and move times, but they do indicate a comfortable margin for achieving the magnitude goal and suggest that a survey based on ten-second integrations in V+R would actually reach somewhat deeper than the V = 24.0 goal.

While it is true that useful survey data can be obtained in several of the color bands, colorimetric data obtained from different filters are of little value in the near-Earth asteroid survey. The reason is, first of all, that most detections will be near the limiting S/N, about 5 or so, and thus the colorimetric precision obtained will be of the order of 0.2 magnitudes, hardly precise enough to be useful. Furthermore, many small asteroids have lightcurve amplitudes of 0.5 magnitude or more, and periods of rotation so short that images in another color, even 15 minutes apart (these small objects tumble quickly!), could be so affected by changing brightness that the inferred color difference might be deceptively wrong. On the other hand, it may well be that the poorer astrometry that will result due to chromatic aberration in a particularly wide filter mitigates against using this for the NEO search, and such broad-band data would be much less useful for the other science goals of the LSST. It is an open question for further study whether it is more efficient to run the NEA survey through standard filters as part of a unified observing program with the other scientific projects described in this document, or to devote a smaller fraction of the total time running an optimized survey for NEAs separate from other programs.

---

[7]Calculated using the 8.4m LSST exposure time calculator, `http://www.ctio.noao.edu/lsst/etc/`

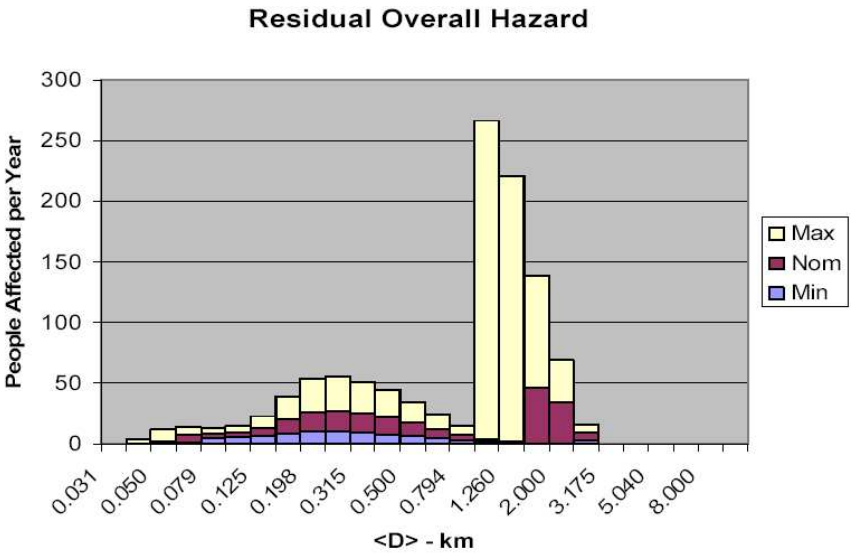

Fig. 3.— The estimated residual hazard from the fraction of the population that will remain undiscovered after the Spaceguard Survey in 2008. Assuming no asteroid on a collision course with Earth is found, most of the risk from large impacts will have been retired, but very little progress will have been made in reducing the risk from smaller impactors. Thus the relative importance of small impactors becomes greater. This is Figure 3-11 in the NASA report.

Table 1. Exposure times required to reach S/N = 5 for a constant size NEA ($V = 24.0$) in $0.7''$ seeing in various bands. These are calculated scaling from actual performance of the CTIO 4-m, and taking the more optimistic numbers from the LSST exposure time calculator.

Color	magnitude	exp,sec. CTIO	exp,sec. 8.4m LSST
B	24.8	21.0	11.5
V	24.0	8.15	4.24
R	23.6	7.50	3.90
I	23.2	10.8	5.92
V+R	(24.0)	5.5	2.9

Returning to the matter of cadence, the plan described in Appendix A would obtain two images 15 minutes apart; this will suffice to identify moving targets (even trans-Neptunian objects with sufficiently accurate relative astrometry) from stationary transients, and at the same time allow unambiguous linkages of images in the two frames. The rates of motion obtained are sufficient to link images of the same objects at least several days later. Thus the repeat time of three days is acceptable for linking objects and obtaining orbits. In fact, up to 5–6 days separation will still allow most objects to be linked, so there is room in the "universal cadence" for interruptions due to weather or equipment problems. Pairs of observations on two nights are a sufficient minimum to obtain a preliminary orbit, although a third night is highly desirable to confirm the linkage and to further improve the orbit. A functional minimum survey strategy is thus to cover the sky area three times during a lunation, which takes a total of 9 observing nights using the cadence of Appendix A. With 50 mas astrometry, an arc of only six days (three observing nights separated by three days each) is sufficient to obtain satisfactory orbits for the majority of objects. The rare few exceptions can be targeted for special follow-up if the preliminary orbit is insufficient to rule out PHA status.

The cadence plan should be regarded as provisional. Much remains to be done to evaluate optimum strategies in terms of sky area coverage, number of repetitions of coverage, and time spacing of repeat coverage. As noted above, coverage of a narrower ecliptic band than 40 degrees might be almost as effective in finding PHAs; the time might better be spent covering a narrower band more often, if only to guard against failed detections near the limit by lightcurve variation. Having three passes a night instead of two could arguably improve the average threshold for successful detections (only two out of three required instead of two out of two) by two or three tenths of a magnitude, well worth the extra time. Optimizing the observing strategy is not urgent beyond that needed to be sure the telescope design is flexible enough to accommodate an optimized strategy when it is determined.

3.5. Summary of Requirements

This chapter has emphasized the study of NEAs from the point of view of the hazard they represent for the Earth. In Chapters 4 and 9 we discuss science that can be learned from studies with LSST of Kuiper Belt Objects and Main Belt Asteroids, respectively. Since NEAs and even PHAs will be detected at distances of about 1 AU or more, their motions will be nearly indistinguishable from MBAs, and in fact a substantial fraction of them will be indistinguishable from MBAs until preliminary orbit solutions are obtained. Thus it will be necessary to catalog and track MBAs if for no other reason than to be able to separate PHAs from this background of "noise". Partly because of this, the requirements that the NEA program puts on LSST are largely shared by both the main belt asteroid and Kuiper belt object programs; these data will also be ideal for studies of comets as well.

With this in mind, the following are the requirements:

Given that the answer to many of these questions are statistical in nature, there are no hard requirements except to optimize the detectability of the faintest possible moving objects.

- *Total area of sky imaged at a time*: As large as possible, all else being equal. A 3° circular area (~ 7 square degrees) is good.
- *Depth and dynamic range of a single exposure*: $R \sim 24$ at $5\,\sigma$ above the noise.
- *Depth and dynamic range needed in stacked exposure*: There is no direct need for this for NEA detection. This does become important, however, in studies of more distant objects. For example, coadding images of moving objects from the same epoch can improve detectability of cometary activity as statistically significant deviation from the local PSF.

One could also use image *subtraction* for detection of image motion in crowded fields.

- *Requirements on seeing, PSF, and pixel size*: Images should be well-sampled (i.e., > 2 pixels across the seeing disk) for optimal astrometry. The better the seeing, the better the astrometry and the depth; see separate requirements for these.

- *Filters*: We have seen that we reach acceptable S/N in the V, R and I filters; a broader-band filter (e.g., $V + R$) would allow us to go deeper. If multiple filters are used, exposures should be adjusted so that a "typical" solar system object reflecting sunlight reaches similar signal to noise ratio in all bands.

- *Photometric accuracy*: In crowded fields of moving objects, brightness may serve as a constraint for link-ups between epochs, so the best possible photometric accuracy is desirable. This of course is mitigated by the rapid variability of asteroids due to tumbling and phase angle effects, especially for the small NEAs.

- *Astrometric accuracy*: The better the astrometry, the better the motion vector is defined, and the better the resulting orbits. This needs quantification, but accuracies of tens of milliarcsec per epoch would be substantially better than the current state of the art for asteroid searches.

- *Requirements on sky darkness and photometricity*: The darker and clearer the sky, the deeper and more accurate the data. Photometry can be done on non-photometric nights by calibrating against previously photometered stars (but work will need to be done to determine the optimal smoothing scale on which to do this comparison). Otherwise, the requirements must be consistent with goals of individual exposure depths and the desire to meet survey goals in of order ten years.

- *Speed of data reduction*: Needs quantification, but fast-moving objects need to be recognized on timescales short enough that they are not lost, especially for any NEAs that may require follow-up on other telescopes.

- *Specialized data analysis tools*: A sophisticated moving object detection algorithm and in-house preliminary orbit linkages. In general, the linkage and orbit determination problem may be quite challenging, given the "noise" of main belt asteroids.

Morrison, D. editor 1992, The Spaceguard Survey Report

Near-Earth Object Science Definition Team 2003, *Study to Determine the Feasibility of Extending the Search for Near-Earth Objects to Smaller Limiting Diameters* (NASA: August 22, 2003)

Schweickart, R.L., Lu, E.T., Hut, P., and Chapman, C.R. 2003, Scientific American, November issue.

Toon, O.B. et al. 1997, in *Near-Earth Objects*, Ann. NY Academy of Sciences, 822, 401

Ward, S.N. & Asphaug, E. 2000, Icarus 145, 64

4. Kuiper Belt Objects

4.1. Kuiper Belt science goals for LSST

The Kuiper Belt (and other distant small body populations) consists of remnants of the early accretion and evolution of the Solar System. In the closer Solar System, runaway and oligarchic growth of solid bodies led to the production of the giant planets, which subsequently ejected most of the remaining planetesimals with perihelia interior to Neptune. In the outer Solar System, however, runaway growth was for some reason truncated, and the Kuiper Belt region still contains a portion of the planetesimal population. Since objects in the 10–1000 km range in extrasolar planetary systems are likely to remain unobservable for many decades, the KBOs represent our only chance to study directly this phase of planetary system formation.

The Kuiper Belt is not dynamically pristine: the eccentricities and inclinations of the known KBOs are substantial, in the sense that accretion could not have occurred in the present dynamical state. Figure 4 illustrates the dynamical state of presently known outer Solar-System bodies. There is a drop in the space density of $\gtrsim 40$ km objects beyond 50 AU that is unexplained. There is a clear correlation of the size distribution with dynamical state, with the largest bodies found exclusively in higher-excitation orbits. These and other current data indicate that the Kuiper Belt contains clues to one or more major events in the history of the outer solar system. The history of accretion, collisional grinding, and perturbation by existing and vanished giant planets is written into the joint distribution of KBOs over orbital elements and size. Colors and compositions of KBOs are clearly diverse and correlated with dynamical state, but the physical origin of this diversity is also unknown. Light curves of KBOs also give information on their shape and surface inhomogeneities; from this one can constrain the angular momentum distribution and internal strengths of the bodies.

A high-throughput telescope such as LSST has the power to discover tens to hundreds of thousands of new KBOs, map their orbital distribution, and determine colors and time variation for many or all of these. The joint distribution over these quantities will allow us to disentangle the history of the outer solar system. The discovery of such a large number of KB objects is desirable for several reasons:

1. Structure in the dynamical (or other joint) distributions becomes apparent only with large numbers of objects, to reduce the shot noise in the phase-space density of KBOs and to find niche populations that likely provide strong clues to the origin and evolution of the belt.

2. Higher object counts arise from more complete sky coverage and/or greater depth. Since the Kuiper Belt has an outer edge, fainter KBOs are smaller KBOs. There is a turnover in the KBO size distribution below ≈ 100 km diameter, presumably since smaller objects are susceptible to collisional disruption and have been ground away since the accretion epoch. Understanding how the erosive turnover depends upon dynamical variables and colors will show when the erosive transition occurred for each dynamical family.

3. More complete sky coverage will ensure the discovery of important but rare objects. With only ~ 1000 KBOs known, we are still discovering objects that force us to revise our basic scenarios (e.g. 2000 CR105, an object with highly elliptical orbit but perihelion beyond reach of Neptune).

4. The upper envelope of rotation rates is an indicator of the physical strength of the bodies, since rapid rotation can cause breakup. This is better defined with larger samples.

5. Having a more complete sample of KBOs provides the grist for mission targeting decisions for future NASA missions to the Kuiper Belt. That is, the more we know about the physical properties of the KBOs, the better we can design missions to the Kuiper Belt to maximize scientific return. Moreover, the larger our

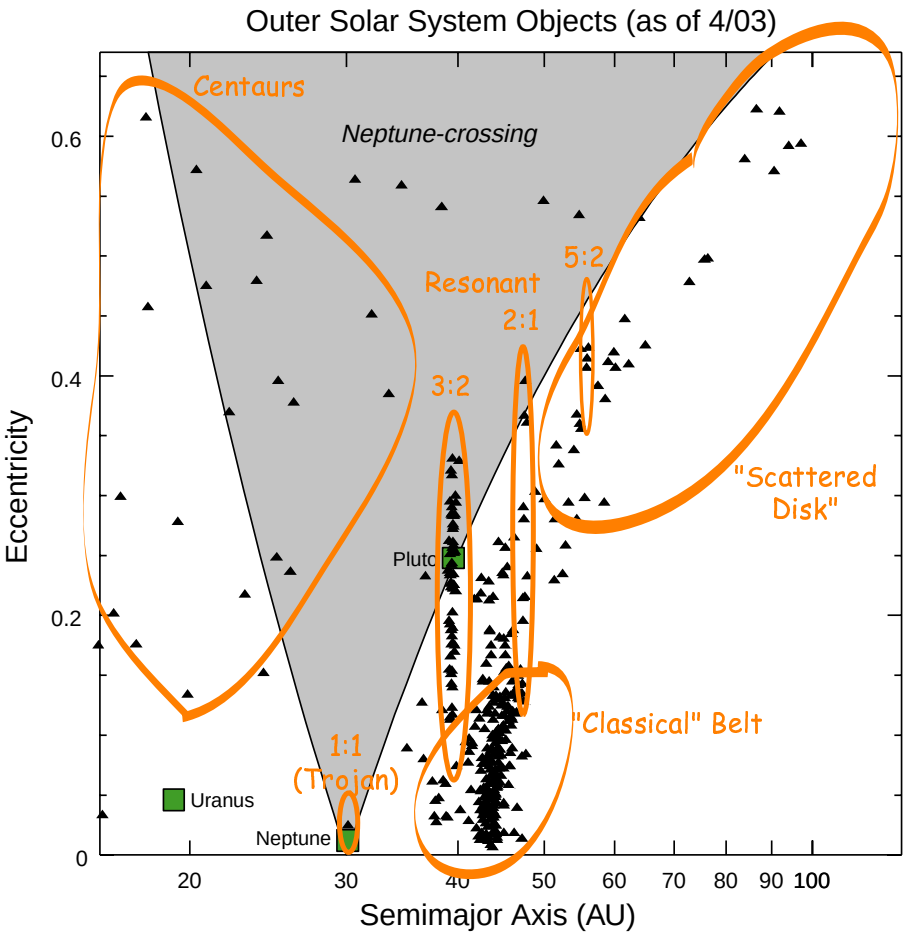

Fig. 4.— The distribution of known outer Solar System bodies in the $a - e$ plane is shown, restricted to those objects with reliable orbital parameters. The readily discernable dynamical families are crudely outlined: the Centaurs have non-resonant orbits that cross the giant planets and are presumably short-lived; the vertical groups are in mean-motion resonances with Neptune (*e.g.* Pluto); the high-e bodies with perihelia near Neptune are the "Scattered disk"; and the lower-e bodies with a between the 3:2 and 2:1 resonances are called the "classical" KBOs because they most resemble the original conception of a pristine residual disk. Even this group, however, has mean e and i much higher than expected for a primordial disk remnant.

sample of KBOs, the more likely we are to find an object with the optimal combination of orbital parameters and physical properties for a visit by a spacecraft.

It should therefore be our goal to use the uniquely high throughput of LSST to increase our knowledge of the KBO population to the extent possible.

Because KBO studies are still in the exploratory era, it is not possible to define a single measurement that must be done and which can be used to produce a quantitative floor on LSST specs and cadences. Nor can we say definitively what number of KBOs with orbits, colors, and/or light curves would be "enough." We can note the following: with the current sample of $\sim$ 800 objects, there are dynamical types that are represented by only one or two instances (e.g. the Neptune Trojan 2001 QR322). Even the most basic correlations between color or size and dynamical properties are marginally detectable. A 100-fold increase in the cataloged population would seem desirable to find sufficient numbers in the known dynamical classes to make meaningful measurements of

size/color trends within such classes. It is also clear that an extension of the well-surveyed population to smaller sizes (fainter limits) is critical to understanding the accretion/erosion history. It is further likely that there are dynamical classes that remain undiscovered in current data. This is emphasized by the recent discovery (Brown et al. 2004) of the $R \approx 22$ object Sedna near its 80 AU perihelion. Sedna may be the sentinel member of a new dynamical class, the inner Oort cloud, with hundreds of members potentially within the reach of LSST.

4.2. Required Signal-to-Noise for KBO Science

The specific requirements for different aspects of KBO science are:

- **Detection:** $S/N > 7$ is required to distinguish TNOs from noise fluctuations, since the objects are very sparse on the sky. This S/N must be acquired in a short time period (see below).
- **(P)Recovery:** $S/N > 5$ is required to recover a known TNO. The threshold for false positives is lower because we do not have to search the entire phase space. The discovery observation does not have to precede the recovery observation, so any detection observation can also serve as recovery.
- **Light Curves:** The amplitude of known TNO light-curve variations range from ≈ 1 mag to zero. A properly phased light-curve will hence require multiple points with S/N of tens to ≈ 100, depending upon the amplitude. A cumulative S/N of 25–200 is thus required to detect the variability, depending upon amplitude. Known light-curve periods are 0.2–1 day, so the observations must span multiple nights, but simulations are needed to determine how well periods can be determined from observations spread over many periods.
- **Color:** KBOs clearly have diverse colors, and they vary by tenths of magnitudes. Hence $S/N \gtrsim 100$ in each of two visible bands is desirable for accurate assignment of KBO colors. Indeed, two bands may be sufficient, as the spectral energy distribution of most KBOs is well-characterized by a single power law. Since KBOs vary on several-hour time scales, observations in different bands must either be within $\ll 1$ hour of each other, or spread over many periods.

The power of an LSST, therefore, is not just in extending the magnitude limit for detection of large numbers of KBOs, but a proper observing cadence can also greatly increase the number of objects observed with sufficient S/N to obtain meaningful colors and light curves. At present, for example, $\lesssim 10\%$ of known TNOs have well-measured colors, and only 1–2% have variability characterized.

4.3. The "Shallow" LSST sample

We will assume for the KBO discussion that there will be a mode of LSST operation centered on NEA detection in which (nearly) the entire visible hemisphere will be imaged in a series of tens of 10–20 second exposures over the course of each year, as described in Appendix A. We will refer to this as the "shallow" survey and to KBOs that can be detected (at 5σ significance) in a single 20s exposure as "bright." Depending on the parameters of the telescope, this will be $R \lesssim 24$, which corresponds crudely to 150 km diameter.

Current data (Figure 5) show that the sky density of $R < 24$ KBOs near the ecliptic plane is $\approx 3\,\mathrm{deg}^{-2}$, roughly equally split between the "classical" Kuiper Belt—a low-eccentricity, low-inclination ($i \lesssim 5°$) population peaked near $a = 42$ AU—and higher-excitation populations, including Neptune-resonant and "scattered" orbits, with a half-width on the sky of perhaps $\approx 20\,\mathrm{deg}$. The total number of "bright" KBOs on the sky is therefore

$\approx 2 \times 10^4$. LSST would easily discover virtually all of these objects and determine high-quality orbits from the shallow survey, since a good orbit will require only 4–6 detections over the course of 2–3 years. This is a roughly 20-fold increase over the number of presently known KBOs.

The job of detecting all the bright KBOs is in fact *too* easy for LSST, in the sense that a telescope with lower étendue will be able to sweep the sky the required 3–4 times to find all the bright KBOs. The CFHT Legacy Survey will conduct a survey of this depth over ≈ 1000 deg^2 centered on the ecliptic in the next six years. The four-telescope incarnation of Pan-STARRS or a similar project will likely have discovered all the bright KBOs by the advent of LSST.

What does the higher throughput of LSST gain us for bright KBOs? LSST will acquire 100 or so observations of each bright KBO over its operative lifetime, as opposed to just a few. This would enable important new science beyond knowing the orbital distribution of the bright KBOs:

(a) Colors: The S/N required for color is well beyond that required for detection; the LSST will give color info for all bright KBOs due to the many repeat visits on the full sky, so the joint color-magnitude-orbital distribution will be known for all bright KBOs. Note that the long-term, random time sampling of the LSST shallow survey will give magnitudes properly averaged over light curves.

A clear requirement is that the NEA survey be split between at least two colors.

(b) Light curves: The 100-or-so observations of each bright KBO can be searched for a light curve period due to rotation of an elongated body, adding amplitude of variation as another variable for which the bright-KBO distribution is fully characterized. Some simulation work is required to test the feasibility of period recovery over such long time scales, and to explore favorable timing schemes. In optimal circumstances, it should be possible to model the light curve for a model of the elongation and orientation of the asteroid. In addition, the data can be searched for coma outbursts and variability due to a binary that goes into eclipse. It seems likely, however, that light curve amplitudes will be measured for many thousands of KBOs, with periods determined for many of them.

4.4. A Deep KBO Survey

We propose here an additional cadence for LSST observations that unleashes the full power of LSST for KBO discovery and study by extending the KBO sample well past the $R < 24$ limit.

Longer integrations are of course necessary to discover fainter KBOs. Near quadrature, KBO apparent motions are $\lesssim 1''$ per hour. A one-hour series of short integrations can be summed to track all such motions, and with an imaging FWHM of $0\farcs5$ or larger, the number of required trial sums is of order 10, which remains in the realm of computational feasibility. We will baseline, then, a survey in which the LSST maintains a pointing for a contiguous hour.

With the effective exposure time increased from 20s to 3600s, the detectable flux (assuming background limit) drops by a factor 13, or 2.8 mag. Only a handful of objects this faint have been detected, but estimates of the sky density suggest this implies a 25-fold increase in the number density of observable KBOs. It also reduces the limiting mass for KBO detection by factor of 50. It is in this mass range that the transition to the putatively erosion-dominated regime occurs, so the collection of large numbers of KBOs in this range will allow comparison of the collisional history of the various dynamical classes. Both the increased number density and the extension to smaller KBO sizes will enormously increase our ability to use the KBOs to diagnose the history of the outer solar system.

The KBO science return will be greatly amplified by an observing mode in which 1-hour segments are devoted to a fixed pointing.

The requirement for useful determination of orbits is likely to be that 3 or 4 detections must be made over a time span ≥ 12 months. A candidate cadence, for example, is:

1. 1 hour at first quadrature year 0.
2. 1 hour at second quadrature, year 0.
3. 1 hour at second quadrature, year 1.

Simulations are needed to determine the trade of visits vs orbital accuracy. The following points about this cadence are clear, however:

- A given 1-hour visit may be done with two or more filters, as long as all filter give good S/N on solar-colored objects. Interlacing filters would give high-accuracy colors for all objects $R \leq 25.5$ within the surveyed area.
- Timing of the KBO visits is not critical; in most cases, delaying a revisit of a field until the next quadrature is not fatal.
- Full sky coverage is not required (indeed not practical—see below) but any partial sampling of the sky should be reasonably uniform in ecliptic longitude, concentrated within 20° of the ecliptic.
- **Visibility of the full ecliptic is an important criterion in the site selection.**

Each LSST field searched for faint KBOs will take a total time investment of 3 hours, and cover 7 deg^2. Taking 170 hours per lunation of dark/grey time, efficiency factors of 0.75, 0.75, and 0.95 for clear skies, good seeing, and uptime, respectively, there are 1200 candidate hours per calendar year of LSST operation. If we presume that a fraction $f_{\rm deep}$ of time is devoted to the deep cadence, then in a 10-year lifetime we can survey 27,000 $f_{\rm deep}$ deg^2 of the sky, or $1.8 f_{\rm deep}$ of the total area of sky within $\pm 20°$ of the ecliptic. If $f_{\rm deep} = 0.1$, then we would expect a total of $\approx 10^5$ detected $R < 27$ KBOs. Roughly 25% of these would have high-precision color determinations.

Note that with these three observations in hand, one can now leverage the accumulated NEA survey data for these fields: the orbit can be tuned to higher precision by fitting to the 1–2 hours of 20-second exposures that have accumulated over 10 years. Simulations are required to determine the magnitude limits to which useful light-curve data could be extracted from the combined deep/shallow database.

A modest investment in long-integration mode for LSST would yield a five-fold increase in detected KBOs and those with useful colors, and push into a different physical regime of KBO sizes. This mode would likely be useful for other domains of time-variable astronomy as well.

4.4.1. Occultation Searches

KBOs of diameter $\lesssim 10$ km are currently undetectable even with HST and searches for their light are likely to remain infeasible until at least the launch of JWST. Known comet nuclei have diameters as small as ~ 1 km, so the Kuiper Belt precursors of these objects have $m \approx 35$, far too faint even for JWST or 30-meter telescopes, yet we would like to extend our knowledge of the KBO size distribution to this size so we can understand the origin of the short-period comets. Such bodies are, however, detectable through occultations: each star near the ecliptic plane will be occulted by a KBO for $\lesssim 1$ second roughly once per thousand years. Hence monitoring many thousands of stars will allow the detection of km-scale KBOs. A pilot occultation survey is about to begin (Taiwanese-American Occultation Survey; C. Alcock, US PI) with a network of 4 small telescopes. The high throughput of the LSST can facilitate a powerful occultation survey, and the large-diameter telescope would allow detection of low-amplitude

occultations caused by sub-km KBOs. The technical requirements are that: **the telescope be operated in the deep-survey mode, staring at a given patch of sky for a significant time before slewing;** and **the focal plane must be equipped with high-speed (CMOS) detectors that can read out the bright-star images at several Hz.** If the CMOS readouts are non-destructive or confined to the bright stars, then the occultation survey can proceed in parallel with the deep KBO survey. The deep-survey mode could also be sensitive to transients on time scales as short as 1 s or as long as 1000 s.

4.5. Required technical specifications

Here we comment on the figures of merit for telescope engineering that are relevant to the KBO science.

- The figure of merit for FOV, aperture, and image quality is the usual point-source quantity (FOV $\times$ $D/\mathrm{FWHM})^2$. In fact it is the time-averaged inverse of this quantity that is relevant—the canonical 1-hour exposure time can be trimmed dynamically in good conditions, if the telescope optics are good enough to take advantage of good seeing.
- Filter choice: for the shallow survey, KBO science prefers that at least part of the survey make use of the filter that optimizes S/N for solar-colored point sources. However to obtain color information, a single wide-band filter is not optimal. Some rotation between g, r, i, and "wide-V" filters is desired.
- Filter choice: for the deep survey, a wide-V filter might provide the best detection limit, perhaps a gain of a factor of 1.5 in object counts. But cycling between narrow filters, e.g. g and r, with i in brighter time, will increase the scientific yield from color information, and may provide a better match with other deep-survey goals.
- Astrometry: KBO orbits will improve usefully as astrometric accuracy improves. A global astrometric frame with errors $\ll 0.1''$ is desirable, though not required.
- Pre-survey: a pre-survey is somewhat useful for KBOs in that it provides a subtraction template for the shallow survey. The deep survey would be too deep for an all-sky pre-survey to serve as subtraction template. The deep survey will have to serve as its own subtraction template.
- Photometric calibration accuracy: 0.02 mag is probably a requirement, better can be used.
- Read time, slew times, overheads: for the shallow survey, the optimization for efficiency is driven by NEO requirements. For the deep mode, exposure times may lengthened and slewing is reduced, so demands on overheads are substantially looser.

4.6. Open Questions

We need further study of the following questions:

- Are 3 (or 4) observations at successive quadratures sufficient to localize the orbit to desired accuracy?
- What kind of tiling strategy maximizes efficiency of a subsampling of sky while minimizing the loss of objects off the FOV over the orbital arc?
- For a deep survey, what kind of filter cadence maximizes the scientific yield for variability studies (cf., § 5) and for the accumulation of a valuable deep static image (cf., § 7)?

- How well can light curve amplitudes and/or phases be recovered from observations taken over a time baseline of hundreds or thousands of periods?
- What is the technical feasibility and event rate for an occultation survey using a CMOS LSST?

Bernstein, G. et al. 2003, preprint (astro-ph/0308467)

Brown, M.E., Trujillo, C., & Rabinowitz, D. 2004, ApJL, in press (astro-ph/0404456)

Durda, D. D. & Stern, S. A. 2000, Icarus, 145, 220

Farinella, P., Davis, D. R., & Stern, S. A. 2000, Protostars and Planets IV, 1255

Luu, J. X. & Jewitt, D. C. 2002, ARA&A, 40, 63

Malhotra, R., Duncan, M. J., & Levison, H. F. 2000, Protostars and Planets IV, 1231

Morbidelli, A., Brown, & Levison 2003, in *First Decadal Review of the Edgeworth-Kuiper Belt*, in press

Tegler, S. C. & Romanishin, W. 2003, Icarus, 161, 181

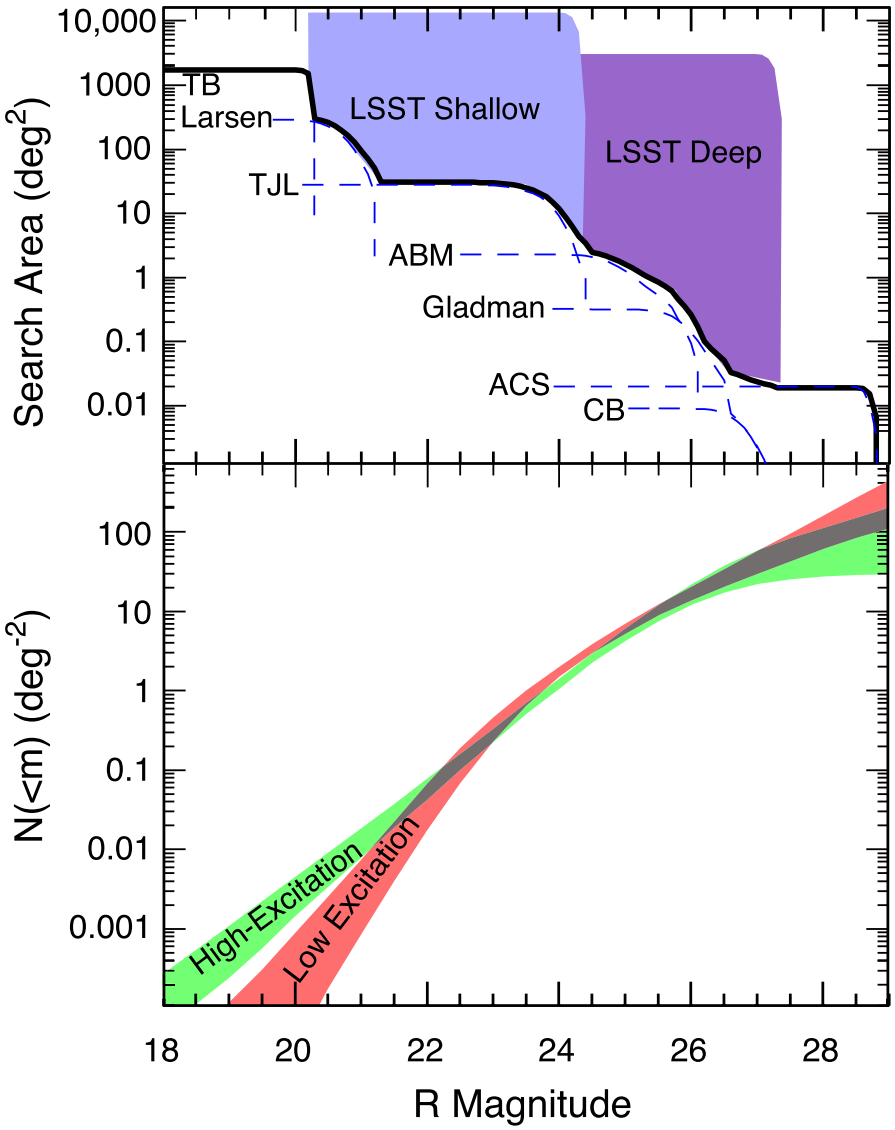

Fig. 5.— The top panel shows the area of sky surveyed to date (restricted to published data within 3° of the invariable plane), and the coverage that should be possible in shallow and deep modes of the LSST over its lifetime. In the lower panel is the cumulative sky density of KBOs (near the invariable plane) vs limiting magnitude, as derived from the top-panel data (Bernstein *et al.* 2003). Dividing the sample into high- and low-excitation KBOs reveals a clear difference in magnitude, and hence size, distribution of the objects—a clear sign of different evolutionary histories that has yet to be fully explained.

5. The Variable Universe: Explosive Transients, the High Energy Sky, and Variable Stars

The LSST, with its large aperture and frequent, repeated observations of a large fraction of the sky, will open a new frontier in time domain astrophysics. A number of surveys over the past decade have explored the optically time variable sky, including those looking for gravitational microlenses and variable stars (MACHO, OGLE, MOA, EROS, AGAPE, PLANET; tens of square degrees sampled nightly or more often roughly to 20th magnitude), moving objects (SpaceWatch, LONEOS; most of the sky covered a few times a month to 19th magnitude), gamma-ray bursts (ROTSE, RAPTOR, large areas of sky going quite shallow), and supernovae and other faint objects (SCP, DLS, going to 25th magnitude or deeper in several bands over quite small areas). The QUEST variability survey (Vivas et al. 2001) uses a 1-m telescope and a large drift-scan CCD camera; it is one of the closest current analogs to what LSST will be capable of in time-domain studies. Relative to QUEST, the LSST will increase the area with extensive multi-epoch observations by a factor of about 30, provide multi-epoch observations to a limit about 4 magnitudes deeper, and extend the short time scale limit from a few hours to a few minutes. Since even two observations, aided by accurate colors, allow for exciting new discoveries, LSST will usher in a new era, likely placing time domain science at a level commensurate with spectral and high-resolution work.

The LSST will allow us to study the optical variability of exotic objects such as accreting neutron star or black hole binaries, blazars – the most extreme end of the AGN phenomenon, and transients such as gamma-ray bursts and X-ray flashes, at least some of which are associated with the demise of massive stars. Most of these phenomena are predominantly identified by their variable high-energy (hard X-ray and gamma-ray) emission. By teaming with future missions such as the *Gamma-Ray Large Area Space Telescope (GLAST)* and the *Energetic X-ray Imaging Survey Telescope (EXIST)*, LSST will provide the multiwavelength temporal monitoring essential in many cases to help unravel the extreme physics taking place in these objects. By searching for the optical afterglows of GRBs, including those with no accompanying gamma-ray flash, LSST can constrain the explosion geometry. Perhaps most exciting is the possibility that LSST will discover new classes of transients, possibly associated with binary mergers. This chapter describes some of the possible science to be addressed in the variable, high-energy and transient sky.

5.1. Gamma-ray Burst Afterglows - Orphans and Light-curves

Catastrophic stellar explosions, such as supernovae and gamma-ray bursts, produce optical transients decaying with timescales ranging from hours to months. Some classes of GRBs (those with gamma-ray durations in excess of about two seconds) result from the explosion of massive stars. These events are known to produce bright optical flashes decaying with hour timescales, resulting from a reverse shock plowing into ejecta from the explosion. ROTSE I detected one such optical flash at a redshift of 1.6, which is the most luminous optical source ever measured ($M_V = -36.4$, Vestrand et al. 2002b). In addition, GRBs produce fainter, longer-lived (decay timescales of days to a week) afterglows associated with collimated (jet-like) relativistic shocks expanding outward into the circumstellar medium. As the outward-going collimated shock evolves, the spectral peak of the afterglow moves from the X-ray to the optical and then radio.

Current observational evidence suggests that most GRB ejecta are strongly collimated, with jet solid angle of $4\pi/\Delta\Omega \geq 100$ (i.e., angles $\leq 10^\circ$ or so). The GRB is thought to be produced in an ultra-relativistic outflow with a Lorentz factor of $\Gamma \sim 100$. As the outflow decelerates by sweeping up the surrounding medium, the relativistic beaming angle $\Theta_j \sim 1/\Gamma$ increases, and the emission is visible from larger off-axis angles. Since the peak frequency of the emission decreases as the jet decelerates and expands, jets that were invisible during the gamma-ray phase (because they were viewed off-axis) when the beaming angles are small will become visible at later times, when the afterglow emission is peaking at optical or radio frequencies. These are called "orphan" afterglows. Thus many

more optical and radio afterglows should be observable than GRBs themselves.

The rate of detectable orphan afterglows depends sensitively on the collimation angles, so that determining the orphan afterglow rate will directly constrain the geometry of the outflow. Determining the shock geometry is essential for deriving both the explosion energetics and the cosmic GRB rate, and in addition would help constrain models of the shock emission.

Accurately constraining the orphan afterglow rate requires repeated observations of the same fields to identify candidates, and a sensitivity to transients to $R \geq 24$. Assuming each field is visited four times in a month, which should be sufficient to establish fading behavior, LSST could detect 500 events/year (Totani & Panaitescu 2002) in its all-sky survey mode, assuming that current theoretical estimates of collimation angles are correct.

5.2. Binary Mergers - Gravitational Astrophysics

A longstanding goal in astrophysics has been to detect mergers of compact objects; neutron star – neutron star or neutron star – black hole binaries whose orbits gradually decay as they lose orbital energy to gravitational radiation. Gravitational wave observatories like LIGO and LISA aim to study the final stages of the inspirals. Mergers were also long believed to be a possible explanation for gamma-ray bursts and X-ray flashes. We now know that long-duration GRBs result from the collapse of massive stars, however it may still be the case that short GRBs may result from inspiral events. In this case, searching for optical counterparts will be a key component in understanding their nature. It may also be the case that mergers result primarily in longer wavelength (optical or radio) transients, with weak or absent high-energy emission.

If LSST succeeds in detecting the optical counterparts to inspiraling compact pairs that produce detectable gravity waves, it can measure their locations in two dimensions, providing an important constraint on the fit to the gravity wave time series. By triggering the acquisition of redshifts using a large aperture spectroscopic telescope such as the Giant Segmented Mirror Telescope (GSMT), these systems can provide important constraints on both our knowledge of gravitational physics and cosmology. The detection of variability with LSST will provide another important way to test the basic nature of gravity.

5.3. The High-Energy Sky - Teaming with GLAST and EXIST

In the upcoming decades the GLAST and EXIST space experiments at GeV and hard X-ray energies, respectively, will observe the high-energy sky with unprecedented sensitivity and with fields of view covering an appreciable fraction of the Celestial Sphere. EXIST will monitor accreting black holes on all scales, from the stellar to the super-massive using a wide FOV (5 ster) coded aperture telescope operating in the 20 – 600 keV hard X-ray band. EXIST will image most of the sky each day, and much of the sky during each 90 minute orbit. GLAST images in the GeV gamma-ray band, covering about 30% of the sky each day, and will primarily detect and monitor high-energy blazars. EXIST and GLAST will obtain light-curves for a large sample of objects, distributed over the entire sky, with temporal sampling on scales from minutes to days (depending on the source brightness).

LSST provides an ideal complement to these missions. EXIST will obtain light-curves for thousands of black holes on timescales of minutes to weeks. By monitoring large regions of the sky, LSST will provide the crucial optical light-curves sampled on day timescales, with baselines extending over years for hundreds of black holes. These multiwavelength variability studies are the best hope for understanding the structure of the powerful jets that emanate from many black holes (Figure 6), the acceleration processes of high-energy particles, and the evolution of the particle energy distribution.

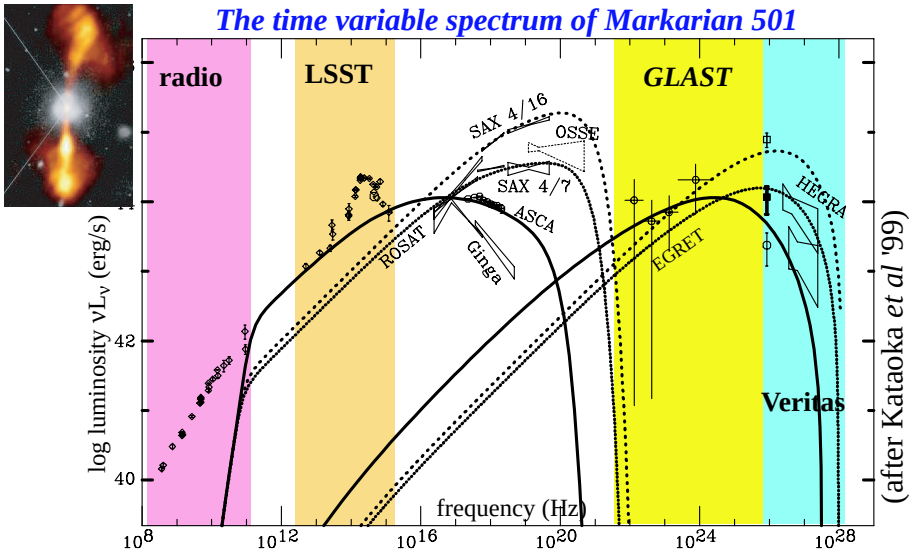

Fig. 6.— Spectrum of the active galaxy Markarian 501. Spectra of blazars typically show two peaks, thought to be synchrotron and inverse Compton radiation. GLAST will observe the upscattered radiation, while LSST observes the primary synchrotron spectrum, and constrains a different part of the relativistic electron spectrum. LSST will obtain photometry in at least two colors several times per month for a large sample of GLAST blazars, crucial for testing physical models of the relativistic jets.

LSST variability information will also be important in helping to find counterparts to unidentified high-latitude high-energy sources. Both EXIST and GLAST will determine the position of sources to an accuracy of 20 – 60 arcseconds, depending on the source brightness. These missions will provide a large catalog of unidentified sources, most of which will be variable AGN at high latitudes (particularly in the GeV band, where most will be blazars). The LSST can identify likely counterparts with its database of long-term lightcurves and colors for the objects in each unidentified source error box.

5.4. Variability and Stellar Astrophysics

Variable stars represent a crucial tool for studying stellar astrophysics. For example, some of the most reliable stellar mass estimates are based on studies of eclipsing binary stars.

The properties of optically faint variable sources are by and large unknown. There are about 10^9 stars brighter than $V = 20$ in the sky, and at least 3% of them are expected to be variable at the few percent level (Eyer 1999). However, the overwhelming majority are not recognized as variables even at the brightest magnitudes: 90% of the variable stars with $V < 12$ remain to be discovered (Paczyński 2000). For example, known RR Lyrae stars are distributed in isolated square patches with the same size and shape as the Schmidt plates used to discover them.

The on-going SDSS variability survey provides a glimpse of the time-domain discoveries that will be made possible by the LSST. Although SDSS is not primarily a multi-epoch survey, overlap between scanlines gives at least two epochs for over 1000 deg^2 of sky to date. These data demonstrate that the optically faint variable population strongly depends on the time scale and also on the object's magnitude. For sufficiently long time scales (a year or longer), variable objects fainter than $r \sim 19$ are dominated by quasars. For time scales shorter than about a month, or at bright magnitudes ($r < 18$), the variable objects are heavily dominated by stars (see below). Indeed, variability is present on all sampled times scales (from a few hours to a few years); a few percent of all

objects are variable at the level of 0.07 magnitude.

Figure 7 shows the SDSS color-magnitude and color-color diagrams for ∼300,000 point sources with repeated SDSS observations (Ivezić et al. 2003a). The overall source distribution is shown as contours, using the mean magnitudes. All objects which changed between two observations by at least 0.075 mag in both g and r bands with a statistical significance of at least 3σ, are connected by lines. The top four panels correspond to observations obtained 733 days apart, and the bottom four panels to observations obtained 3 hours apart. On a ∼2 year time scale, the majority of objects (77%) have colors typical of low-redshift quasars ($u - g < 0.6$). Blue stars consistent with the halo turn off, both in colors and in apparent magnitude, make up 40% of the remaining objects (this subsample may also contain some quasars), and the reddish stars with $u - g > 1.3$ another 40%. RR Lyrae stars, which are recognized by their distinctive SDSS colors (Ivezić et al. 2000, 2003b), comprise 20% of detected variable stars. At the three hour time scale, only about 2% of variable objects have quasar colors. RR Lyrae now make up 35% of the sample, mainly because the number of red stars is much smaller (the surface density of selected RR Lyrae is about 1 per deg^2, nearly independent of time scale and galactic coordinates). The smaller number of red stars is consistent with them being long-period (Mira) variables which don't appreciably vary on time scales of several hours.

The repeated LSST observations will enable time domain science on a scale far beyond the current state of the art. Since even two observations, aided by accurate colors, allow for exciting new discoveries as demonstrated by SDSS, LSST will usher a new era of massive and accurate variability studies, and may place time domain science at a level commensurate with spectral and high-resolution work.

5.5. LSST requirements

GRB Afterglows. Constraining the orphan afterglow rate requires clearly identifying at least several hundred fading afterglows. A key challenge will be separating them from other variable sources. This will require three or four photometric observations during the fading to establish the characteristic $F \propto t^{-\alpha}$ behavior. The LSST will naturally give a large number of observations well before and well after the afterglow; the absence of light from the object will clearly establish the transient nature of the source. Stacked LSST images of the field will go sufficiently deep to identify the host galaxy of the afterglow.

On average, optical transients should be detectable above their host galaxy emission for about 20 – 30 days, when they reach $R \sim 24 - 25$. The characteristic power-law light curve identifies these as GRB afterglows. Color information is unlikely to be useful for filtering background transients, since although afterglows are intrinsically blue, host extinction can make them appear red. The universal cadence described in Appendix A will be adequate to constrain the orphan afterglow rate, assuming that each of the ~ 4000 deg^2 imaged on a given night is revisited on two – three day intervals for at least ten days.

Each year there are approximately 1000 GRBs all-sky. If each LSST field reaches to $R = 24$, each field is covered once every three days, and current estimates of the average beaming angle are correct, twenty times as many orphan afterglows as GRBs will be detectable, meaning that as many as 100 afterglows per month could be discovered. Photometry need only be accurate enough to constrain the fading behavior – ie. to about a tenth of a magnitude (although much more accurate photometry might make apparent more subtle features in the light curve). There is no constraint on the length of individual exposures.

Binary mergers. No theory provides definite predictions for the brightness or timescale of optical transients associated with binary mergers. Relevant timescales for fading afterglows could range from minutes to days. The universal cadence described in Appendix A provides good coverage of phase space for discovering mergers for afterglow timescales of 15 minutes or longer.

To cover short timescales of a minute or less requires three or more exposures in rapid succession so that cosmic ray hits are not mistaken for transients, and a fading source can be established. In addition, these should all be in the same filter, so that color and variability information can be separated. For an instantaneous field of view of f ster, the relation between the all-sky rate of short timescale events a and the rate l at which they are discovered by LSST is:

$$\epsilon a = l \times \frac{4\,\pi}{f}, \tag{1}$$

where ϵ is the fraction of the year that that LSST observes. For example, if $\epsilon = 0.3$ and $f = 7\,\mathrm{deg}^2$, an all-sky rate of $a = 450\,\mathrm{day}^{-1}$ allows LSST to discover ten short events per year. Table 2 shows that if short transients associated with NS-NS, BH-NS, or BH-HE mergers were bright enough to be visible to high redshift, it is plausible that LSST could detect them even if their timescales are as short as a minute. LSST will be much more sensitive if the transient timescales are 15 minutes or longer, since ~ 650 fields of view will be observed twice at 15 minute (or greater) separation each night, so that events with an all-sky rate two – three orders of magnitude smaller can be discovered.

This work requires photometric accuracy of 0.1 magnitude or better. Good astrometric accuracy relative to other objects in the field will aid in identifying associated host galaxies. Rapid data reduction and transient identification (i.e., within 24-48 hours) is required so that followup spectroscopy and observations at other wavelengths can be carried out in order to aid in determining the nature of any detected transients. The followup may require dedicated followup spectroscopy of a large number of objects, as well as light-curve monitoring over periods of days. The latter may be accomplished with auxiliary facilities, and is probably not efficiently done by LSST itself.

High-energy sky. At high Galactic latitudes, extragalactic sources are the best candidates for synergistic observations with GLAST and EXIST. For variable AGN, EXIST should detect 50 – 100 sources bright enough for daily variability monitoring, and GLAST will detect about 20 (5 of which will be bright enough to measure variability on hour timescales). The ideal cadence for studying the brighter objects is to cover about 10 – 20% of the sky with short exposures (most optical counterparts for the variable AGN will be $R = 21$ or brighter), repeating each field on several hour timescales. For identifying counterparts to unknown sources, variability measurements on day to week timescales are relevant, since the goal is to identify the level of long-term variability rather than to cross-correlate lightcurves. Photometric accuracy to 0.1 magnitude is adequate, and there are no unusually stringent astrometric requirements.

Stellar variability. Here accurate and robust photometry is paramount. Previous stellar variability studies have been very much driven by photometric precision and control of systematics and non-Gaussian outliers (cf., the tremendous increase in numbers of gravitational microlensing events found when image subtraction techniques improved photometric precision by a factor of two; Wozniak et al. 2001). This leads to a requirement of photometric precision of 0.02 mag rms; we may well find ourselves wanting another factor of two precision. The error distribution on the photometric calibration should be well-described by a Gaussian distribution; non-Gaussian tails can play havoc with searches for rare types of variability.

merger type	rate in Universe (day^{-1})	reference
NS-NS	800	Burgay et al. (2003)
BH-NS	450	Fryer et al. (1999)
BH-WD	20	Fryer et al. (1999)
BH-He	1000	Fryer et al. (1999)

Table 2: Published estimates of merger event rates.

Searches for variability require a dense covering of sampled time scales. The cadence strategy described in Appendix A may be adequate for these purposes.

Burgay, M. et al. 2003, Nature, 426, 531
Eyer, L. 1999, Baltic Astronomy, 8, 321
Fryer, C.L. et al. 1999, ApJ, 520, 650
Ivezić, Ž. et al. 2000, AJ, 120, 963
Ivezić, Ž. et al. 2003a, *Variability Studies with SDSS*, astro-ph/0301400
Ivezić, Ž., Vivas, A.K., Lupton, H.R., & Zinn, R. 2003b, astro-ph/0310565
Paczyński, B. 2000, PASP, 112, 1281
Totani, T. & Panaitescu,A. 2002, ApJ, 576, 120
Vestrand, W.T. et al. 2002, SPIE, 4845, 126
Vivas, A.K., et al. 2001, ApJ, 554, L33
Woźniak, P.R., *et al.* 2002, Acta Astronomica, 52, 129

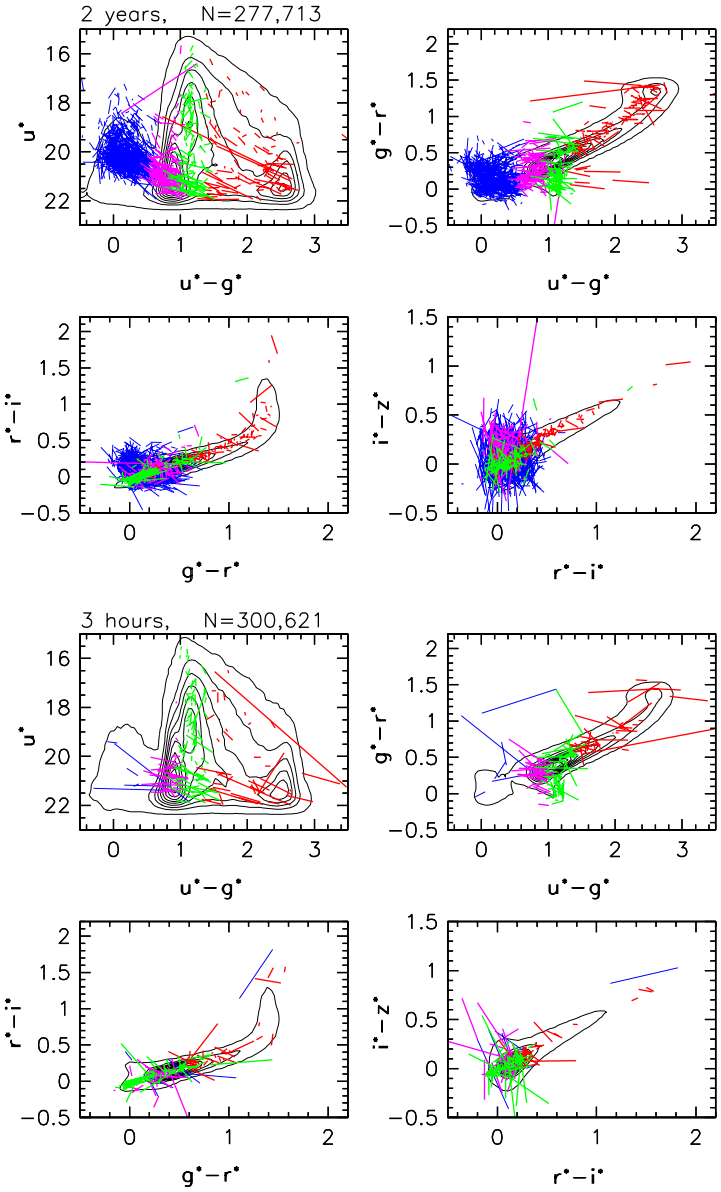

Fig. 7.— The top four panels show variable sources discovered in 75 deg^2 of sky observed twice 733 days apart (Ivezić et al. 2003a). The overall source distribution in SDSS color-magnitude and color-color diagrams is shown as contours, using the mean magnitudes. The two measurements for variable objects are connected by lines, color-coded according to $u-g$ color. The bottom four panels show variable sources discovered in observations obtained 3 hours apart. Note the absence of variable quasars (blue lines, $u-g<0.6$) in the bottom four panels.

6. Stellar Populations and the Structure of the Milky Way

6.1. Introduction

Since Baade's development $\sim$ 60 years ago of the concept of two populations of stars, the study of stellar populations has been essential to our current understanding of galactic structure, evolution, and formation. Measurements of stellar photometry, variability, and astrometry played absolutely crucial roles in the discoveries that the Sun lies many kiloparsecs from the center of the Milky Way, that the Milky Way is one galaxy among many, and that galaxies formed through gravitational collapse. Recent developments based on accurate wide-field stellar photometry and astrometry have made it clear that the Milky Way halo is a complex and dynamical structure, continuing to be shaped by the infall and tidal disruption of neighboring dwarf galaxies. The LSST will allow the application of these techniques on a truly grand scale, laying the path to discoveries of equal magnitude. In this section, we explore the impact that the LSST will have on the study of stellar populations through a number of example cases. We focus on two questions that we particularly think will lead to great discoveries:

- What is the structure and accretion history of the Milky Way?
- What are the fundamental properties of all stars within 200 pc of the Sun?

6.2. The Structure and accretion history of the Milky Way

Over the past decade, both theory and observations have converged on a basic framework for the formation of galaxies. The spectacular measurements of the cosmic microwave background made by COBE and WMAP, the images of young galaxies viewed at high redshift, and the discovery of the cosmic web of large-scale structure by galaxy redshift surveys — all of which are reproduced, to a degree, by numerical simulations — paint a broad picture of the development of mature galaxies from the seeds planted by the Big Bang. A large number of important questions remain unanswered, however, including: 1) How are the earliest generations of stars distributed? 2) What are the accretion histories of galaxies? 3) What is the relationship between the luminous and dark matter in galaxies? There have been a variety of discrepancies between observations and dark matter models on scales of individual galaxies, in particular the concentration of dark matter at the core of a halo and the number of companion galaxies to a luminous galaxy like the Milky Way. A more complete observational picture of the structure of the Milky Way and its environs can go a long way towards clarifying these discrepancies, leading towards new ideas for the nature of dark matter and the formation of galaxies in the early universe. Accurate deep multi-epoch photometry and astrometry over wide fields such as the LSST will be able to provide, will yield insights into these questions.

6.2.1. Tracing the Luminous Halo

Studies of the distribution of F turn-off stars and RR Lyrae stars show that the Galaxy has a stellar halo of steeply falling density ($\rho \propto r^{-3}$ out to 70 kpc; Ivezić et al. 2000; Yanny et al. 2000). RR Lyraes in particular have $M_R \sim 0.5$, and can be recognized to 70 kpc with SDSS photometry. The LSST will reach $m_R \sim 25$ with S/N ~ 50 and $m_R \sim 26.5$ with S/N $= 10$ over the lifetime of the survey; LSST will thus be able to use the numerous main sequence turnoff (MSTO) halo stars, which are characteristically blue and have $M_R \sim 4$, to detect the structure and extent of the halo to distances of $\sim 160 - 300$ kpc. The power of this technique of isolating populations by color in extremely large datasets is being actively demonstrated by SDSS; Figure 8 shows the dissection of the disk, thick disk, and halo components of the Galaxy in a portion of the SDSS survey (Ivezić et al. 2003, in preparation).

With sufficiently accurate colors (1-2%), one can similarly separate giants from main-sequence stars in the halo (Helmi et al. 2003).

The wide field and repeat observations of LSST will make proper motions another powerful tool for the study of the Milky Way halo. While establishing the delivered astrometric accuracy of LSST will require careful modeling and analysis, we imagine that a reference frame of > 100 color- or variability-selected QSOs per LSST field will make it possible to measure positions to the limits set by photon noise and pixel size (cf., the discussion in § 6.3.1). The astrometric error induced by photon noise is characterized by the relation $\sigma = \mathrm{FWHM}/(2 \times SNR)$; using the DRM design and considering MSTO stars with $M_R = 4$, we find that $\sigma \sim (\mathrm{FWHM}/0\farcs5)/(400 \times (\mathrm{d}/20\,\mathrm{kpc})^{-1.65})$ arcsec. This corresponds to a transverse motion $v_t = 170 \times (\mathrm{FWHM}/0\farcs5)(\mathrm{d}/100\,\mathrm{kpc})^{2.65}(\mathrm{N}/100)^{-1/2} \times (\mathrm{t}/10\,\mathrm{yrs})^{-1}$ km/s, where N is the number of visits and t is the timespan of the survey. The expected transverse motion accuracy for halo MSTO stars at 100 kpc over the survey lifetime is thus of the order the velocity dispersion of the halo, ~ 150 km/s, making detailed halo kinematics for distances smaller than 100 kpc possible. RR Lyrae stars, which are brighter than MSTO stars, will have proper motions more accurate by a factor of ~ 14 for the same distance. Proper motions will also be extremely useful for removing contaminants from halo samples, such as foreground white dwarfs.

GAIA[8] is a complementary astrometric space mission being planned by the European Space Agency; it will not go nearly as deep as will LSST. GAIA will determine proper motions to an accuracy of 1 km s^{-1} at 10 kpc, but its magnitude limit means that it will see RR Lyrae stars between 10 and 80 kpc; only rare giant branch tip stars will be visible to larger distances.

LSST will also use variability to isolate halo populations. LSST can detect RR Lyrae with $m_r \sim 23.5$, allowing exploration of halo structure out to 400 kpc, or halfway to M31, with resolution of $\sim 6\%$ in distance. As is discussed in § 6.2.3, one can similarly select classical novae to similar distances, and use them to map the structure of the halo.

What will LSST accomplish? First, it will map the three-dimensional shape and extent of the halo. The tracing of the halo will be done in a highly statistical manner by matching models with specific properties to the number counts and proper motion distributions of nearby halo stars as a function of color and variability properties. The nearly full-sky coverage is absolutely critical for this goal, as narrow slices of sky can easily lead fits of global models astray. Figure 9 (Ivezić et al. 2003) demonstrates this danger of deriving global halo properties from small-area surveys. If only a segment of the SDSS survey were used, one would derive very different halo profiles.

Second, LSST will provide a massive database with which to study the accretion history of the halo. Current studies, performed with much more limited datasets, argue that the nearby halo is actually very smooth (Lemon et al. 2003, Gould 2003). If the halo is composed of streams, then they must number at least in the thousands, with most streams containing only a small fraction of the total number of stars (Gould 2003); a massive sample of stars is thus needed to characterize the streams, if they exist. To best trace the streamers we need both distances and velocities to halo stars. LSST provides the best avenues for obtaining each of these over large areas of the sky. First, RR Lyrae and MSTO stars will be identified throughout the halo and provide distances for any overdensity of stars identified as a potential streamer. Second, proper motion measurements (which will be available for *every* star brighter than $R \approx 24$ in the LSST sky coverage) will provide information on the kinematic coherence of any overdensity. With the tremendous number of stars involved, we will be able both to identify individual streams, and quantify the "lumpiness" of the velocity distribution to constrain the number of streams contributing to the halo population.

[8] `http://astro.estec.esa.nl/GAIA/presentations.html`

6.2.2. Tracing the Dark Halo

LSST can play a role in the study of Galactic dark matter in at least three ways: 1) it will provide dynamical information to map the inner halo via proper motions, 2) it will identify RR Lyrae stars at large distances to greatly increase the number of dynamical tracers in the outer halo, 3) it will provide the most comprehensive survey of the luminous Galactic components.

The distribution of speeds of bound stars in the Galaxy is a direct measure of the depth of the potential; in particular, the upper envelope of that distribution is a compelling measure of the total mass of the halo. From proper motions and photometric parallaxes, LSST will be able to map out this distribution function as a function of three-dimensional position in the halo in detail, leading to a mass model for the Milky Way as a whole far more sophisticated than exists today.

A self-consistent mass model would not only need to explain the velocity distribution function but would need to tie that into the density of objects at larger radii (for example, if we find that many local stars have apogalacticons of 100 kpc, there better be a corresponding population of stars at 100 kpc). Tracing a population in this manner is much more constraining on a model than simply using a satellite galaxy (or two) at large radii. Existing studies (e.g., Kochanek 1996; Wilkinson and Evans 1999; Clewley et al. 2002) have been limited to samples of no more than 100 stars, although recently SDSS data have been used to produce larger samples (Sirko et al. 2003). The LSST will dwarf these studies by many orders of magnitude.

Another probe of the dark halo is via gravitational microlensing towards the Magellanic Clouds, which probes the density of objects of masses within a few orders of magnitude of a solar mass. The LSST can provide critical information on the radial distribution of stellar lensing populations and the tangential velocities of lensing sources. Constraints on halo models from microlensing observations are assumption-laden; for example, Geza, Evans, & Gates (1998) found that brown dwarfs fit the lensing data if the halo is dynamically young and lumpy. With a proper model of the dynamical state of the halo and a full survey of luminous halo tracers such as LSST will provide, one will be able to interpret the microlensing results to constrain the properties of the lenses themselves.

6.2.3. Intergalactic Tramp Stars and Classical Novae

Galactic cannibalism and harassment is a process critical to understanding the evolution of galaxies. Over a Hubble time, many (perhaps most) galaxies suffer one or more close encounters with other galaxies. Tidal tails are often observed in colliding galaxies. Simulations reproduce these features amazingly well, and demonstrate that many stars are liberated from galaxies during collisions. These "intergalactic tramps" have been found observationally as red giants and planetary nebulae in the Virgo and Fornax clusters of galaxies. Our knowledge of tramps outside these systems is essentially nil. Are they common or rare? Novae and LSST offer the opportunity to directly determine the number ratio of stars inside galaxies to intergalactic tramps.

Roughly once per decade a Galactic classical nova attains naked eye brightness. For a few days or weeks the object shines with $10^5\,L_\odot$ before fading back to 15-20th magnitude obscurity. The Milky Way is host to roughly 20 classical novae every year, (as is M31), though only a few are close and bright enough to be detected... often by amateurs.

The physical processes underlying classical nova explosions are extremely well understood. A white dwarf accretes hydrogen-rich matter from a main sequence companion, developing an electron degenerate envelope of roughly $10^{-5}\,M_\odot$ and 1 km in depth. The pressure at the base of the hydrogen-rich envelope eventually becomes large enough to initiate nuclear fusion. This turns into a thermonuclear runaway because pressure does not rise until temperatures in the electron degenerate matter exceed 10^8 Kelvin. The resulting rapid nuclear energy release,

visual luminosity rise and envelope ejection produces a classical nova.

Novae, it turns out, are excellent standard candles. There is a tight and (theoretically) well understood relationship between observed nova peak luminosity and time to decline by 2 or 3 magnitudes from that peak brightness (Shara 1981 ApJ 243, 926). This absolute magnitude-decline time relationship (which is fundamentally due to the tight mass-radius relationship for all white dwarfs) has remarkably small scatter and is independent of underlying binary population metallicity. Novae are seen in very old populations (e.g., in the giant elliptical galaxy M87 and in the globular cluster M80) so they are clearly long-lived. For all these reasons novae are potentially superb tracers of intergalactic tramp stars.

About half of all novae are bright enough ($M < -7$) for long enough (1-2 weeks) to be detectable by LSST out to distances of about $(m-M) \sim 31$, i.e., out to the Virgo or Fornax clusters of galaxies. There are roughly one hundred galaxies with masses comparable to the Milky Way or M31 out to this distance, each displaying about 20 novae annually...roughly 2,000 novae/year in galaxies accessible to LSST. If a (very) conservative 10% of all stars out to $(m - M) \sim 31$ have been ripped from these galaxies, then roughly 200 intergalactic tramp novae will be seen every year by LSST. The ratio of tramp to Galactic novae should mirror that of tramp to Galactic stars. The ~ 1000 tramp novae detected during a 5 year LSST survey will have well determined distances and thus will act as probes of the spatial distribution of all intergalactic tramp stars.

While an ideal observing campaign would image the same piece of sky every night to catch every extragalactic nova at its peak brightness, observations every second or third night will still yield light curves complete enough for very good distance determinations. Novae typically display $B - V \sim 0$ near maximum light, and the light curves in these two passbands are particularly well calibrated, so there is a modest preference for these filters.

A few-times-per-week all-sky survey to 24th magnitude will yield most of the halo (and intergalactic) RR Lyrae stars as far away as M31; and moderate and high galactic latitude contact binaries as far as the Magellanic Clouds. These two populations' distinct light curves will allow their easy identification with the collection of ~ 100 or more observations (to unambiguously identify periods for the variables). Distances are readily determined for contact binaries as there is a tight period-luminosity relationship. The spatial distributions of these stars will be strong constraints on the collision/stripping/cannibalism history of the Local Group.

6.3. A Complete Sample of Stars within 200 parsecs

The solar neighborhood is a unique laboratory for studying important topics such as the stellar and sub-stellar initial mass function and the age of the Galactic disk as derived from the white dwarf cooling curve. In the past, the selection of samples for solar neighborhood studies have had to use magnitude, color, or proper motion as a surrogate for distance, each of which is biased in its own way. By measuring accurate positions for an estimated 10^{10} stars several times per year, the LSST will produce a massive database from which nearby star samples may be selected solely on the basis of their parallax. The LSST will thus provide the first distance-limited catalog of nearby stars subject only to apparent magnitude and celestial coverage limitations. In this section, we describe the measurement of parallax and its uses in more detail.

6.3.1. Astrometry

Although not explicitly discussed by the Decadal Review, the ability of LSST to provide accurate positions for objects is a necessity. Astrometric accuracy at the level of a few tenths of the size of the PSF is part of the basic processing pipeline, but if we reach accuracies approaching those predicted by photon statistics (as limited by seeing, detector quantization, etc.), then exciting and fundamentally important scientific results can

be obtained. We expect that by carefully establishing a reference grid of QSOs, the astrometry will be limited solely by photon noise. For objects with S/N > 100 ($R < 20.5$) in single exposures, the pixel scale will establish a minimum astrometric error of 2 mas per observation, with 0.2 mas achieved by the end of the survey; LSST will thus measure parallaxes with 10% accuracy within a 500 pc volume. For comparison, GAIA, a space-based astrometric mission currently in the planning stages, will measure positions at a given epoch with 1 mas precision at $g = 17$, rising to 17 mas at $g = 22$; by the end of the mission, GAIA will have positions accurate to 25μas at $g = 17$ and 100μas at $g \sim 19.5$. LSST will extend high-accuracy astrometry to stars fainter than $R \sim 20$, where GAIA gives no useful data; at a distance of 500 pc, this magnitude limit corresponds to stars with masses lower than 0.15 $M_\odot$ (Girardi et al. 2000), approaching the hydrogen-burning limit.

6.3.2. The Parallax Survey

LSST offers the opportunity to measure the parallax of every object in its field of view, and thus obtain a complete catalog of hydrogen-burning stars to 500 pc, and of brown dwarfs to tens of parsecs. Although this sounds both obvious and trivial, this project was beyond the reach of the astronomers until the Hipparcos satellite (Perryman et al. 1977) provided these data for the bright stars. Many, most notably Murray (1986) attempted the measurement of the parallaxes for a large number of faint stars using Schmidt telescope survey plates and automated plate measuring machines, but relatively large errors were encountered. Indeed, the Yale Parallax catalog (van Altena 1995) lists values for only 8,112 stars.

The LSST survey of the nearest stars (within 10 pc) will be of great importance to a large segment of the community, and should be available after the first year of operation. If LSST can observe the available sky twice per lunation, then the parallax and proper motion can be separated after a few months, and the large signal from a nearby star can be detected. Of course, more observations will improve the determination and weed out the occasional binary system that fits the detection scheme, but parallaxes with enough accuracy (10% or better) for astrophysical follow-ups will be available in almost real time. It is proposed to make this list available as quickly as possible so that further studies can be started before the stars leave the evening sky. No more will the astrometric studies of newly identified nearby stars (e.g., Dahn et al. 2002 and Vrba et al. 2003) await the curious process of photometric discovery. Astrometry will be used to select the nearby stars without any constraints on photometry or spectroscopy. The discovery of L-dwarfs, T-dwarfs, and other objects with low luminosity are just previews of the new types of objects waiting to be discovered. For example, the cooling evolution of very old white dwarfs is a unique scientific opportunity for LSST. The volume and depth of this survey will discover rare degenerate residues of non-standard stellar evolution.

6.3.3. The Wiggle Survey

Binary stars offer one of the few opportunities to measure the mass of a star, but directed and global searches for binaries are rarely done. Usually, binaries are discovered by accident (corrupting an astrometric or spectroscopic analysis) or through directed surveys looking for companions around specific types of stars (e.g., planets around solar-type stars). Once the binary nature of a system is demonstrated, many further observations are needed if an orbital solution is to be derived. The Washington Double Star Catalog (WDS; Mason et al. 2003) shows 5 systems with semi-major axes larger than 0.7 arcsec and periods less than 20 years. Indeed, the WDS lists orbits for fewer than 2000 systems.

LSST will change all of this. The residuals from the fits for position, proper motion, and parallax will be searched for the signature of Keplerian motion. For systems whose orbital period is shorter than about a quarter

of the available epoch difference, the periodic nature of orbital motion can be sensed and parameterized. Once adequate coverage is available, an orbit can be computed. Modern Fourier techniques can do this automatically on all stars in the LSST archive. A very exciting aspect of LSST is that the astrometry will be done in each of the survey passbands, and the amplitude of the astrometric perturbation can be measured as a function of color. This will allow a statistical identification of the components involved, and allow the list of all binaries to be culled so that the important ones can be handed to the user community for additional observations.

Perhaps the greatest contribution of the Wiggle survey will be in combination with the Parallax survey for the understanding of the star formation process. The LSST data will provide an almost complete inventory of systems within a nearby volume, and the wiggles and common proper motion data can be used to identify the major components of each system.

6.4. Summary of requirements

Here we tabulate the requirements on the system for the Galactic halo, intergalactic tramp star, and solar neighborhood programs.

- *Total area of sky imaged*: The entire visible sky available at any given time observed once per month per filter. Over the course of the survey, this will yield 100 observations of a given area of sky per filter. The intergalactic tramp star program has more stringent requirements: at least every third night to get good light curves, and as much sky coverage as possible. Indeed, a smaller area nightly is better than a large area with sparse coverage.

- *Depth and dynamic range of a single exposure*: $24 < r < 15.5$; the ability to extend this dynamic range to $r = 10$ would be ideal. LSST is the first instrument whose dynamic range and cadence allow it to observe all of the astrometric standards that define the International Celestial Reference Frame (Hipparcos stars and radio QSOs). The LSST catalog will be the pointing catalog for most future ground- and space-based missions, and it is critical that the astrometric reduction have rigid ties at its bright and faint ends.

- *Depth and dynamic range needed in stacked exposure*: $26 < r < 15.5$.

- *Length of individual exposures*: No direct constraints.

- *Requirements on slew time*: Settling time should not impact image quality

- *Requirements on seeing, PSF, and pixel size*: Images should be well-sampled (i.e., > 2 pixels across the seeing disk). The better the seeing, the better the astrometry, the photometry, and the depth; see separate requirements on all of these. The more uniform the PSF across the field, the easier it will be to reach requirements on astrometric and photometric precision.

- *Filters*: The Galactic halo program ideally would use *ugriz* once per month to go deep; *gri* would be acceptable. The solar neighborhood program would need r and i exposures, once per month (but for the first year, twice per month would be preferred, to obtain good initial estimates of parallaxes).

- *Photometric accuracy*: The Galactic halo program needs rms errors of 0.02 mag in photometric calibration; it is possible that even better photometry (0.01 mag) could be used to separate main sequence and giant stars. This goal also requires that the tails of the error distribution be reasonably Gaussian to minimize the false positive rate. The solar neighborhood and novae programs are not as sensitive to calibration errors.

- *Astrometric accuracy*: Relative astrometry should be photon-noise limited, and should reach 2 mas rms per coordinate for S/N > 100. No more than 10% of objects should lie beyond $3\,\sigma$ in the Gaussian distribution. Systematic large-scale astrometric errors (i.e., problems in absolute astrometry) will limit ability to search for coherent flows in the halo; this needs quantifying.

- *Requirements on sky darkness and photometricity*: The darker and clearer the sky, the deeper and more accurate the data. Photometry can be done on non-photometric nights by calibrating against previously photometered stars (but work will need to be done to determine the optimal smoothing scale on which to do this comparison). Otherwise, the requirements must be consistent with goals of individual exposure depths and the desire to meet survey goals in of order ten years.

- *Speed of data reduction*: We require that the list of all stars within 10 pc be available by the end of the first year of observing. The novae program should give list of potential novae within 24 hours of observations.

- *Auxiliary data needed*: The nova program needs follow-up low-resolution (5Å) spectroscopic observations soon after discovery for confirmation of their nature.

- *Specialized data analysis tools*: Specialized astrometric analysis may be needed to get photon noise-limited relative astrometric accuracy. An astrometric reduction package will be needed to solve for parallax, proper motion, and astrometric perturbations.

- *Scheduling of Observations:* Due to the wide passbands and the desire to observe as much of the sky as possible, unmodeled differential color refraction (DCR) will be a major portion of the astrometric error budget. Minimization of the zenith distance for the observations is desirable as is breaking degeneracies between zenith distance and the effects of proper motion and parallax. The modeling of DCR requires multi-band photometric data, too.

Bullock, J.S., Kravstov, A.V., & Weinberg, D.H. 2001, ApJ, 548, 33
Clewley, L., et al. 2000, MNRAS, 337, 87
Dahn, C.C. et al. 2002, AJ 124, 1170
Geza, G., Evans, N.W., and Gates, E.I. 1998, ApJL, 502, 29
Helmi, A. et al. 2003, ApJ, 586, 195
Ivezić, Ž. et al. 2000, AJ, 120, 963
Ivezić, Ž. et al. 2003, astro-ph/0309075
Kochanek, C.S. 1996, ApJ, 457, 228
Mason, B.D., Wycoff, G.L., Hartkopf, W.I. 2003 http://ad.usno.navy.mil/wds/
Morrison, H. et al. 2001 astro-ph/0111097
Murray, C.A., Corben, P.M., Argyle, R.W. 1986, in Astrometric Techniques, ed. H. Eichhorn and R. Leacock (Dordrecht, Reidel), p213
Perryman, M.A.C. et al. 1997, The Hipparcos and Tycho Catalogs (Noordwijk, European Space Agency SP-1200)
Sirko, E. et al. 2003, ApJ, submitted
van Altena, W.F., Lee, J.T., Hoffleit, E.D. 1995 The General Catalog of Trigonometric Catalogs (New Haven, Yale)
Vrba, F.J. et al. 2003 ApJ, submitted
Wilkinson, M.I., and Evans, N.W. 1999, MNRAS, 310, 645
Yanny, B. et al. 2000, ApJ, 540, 825

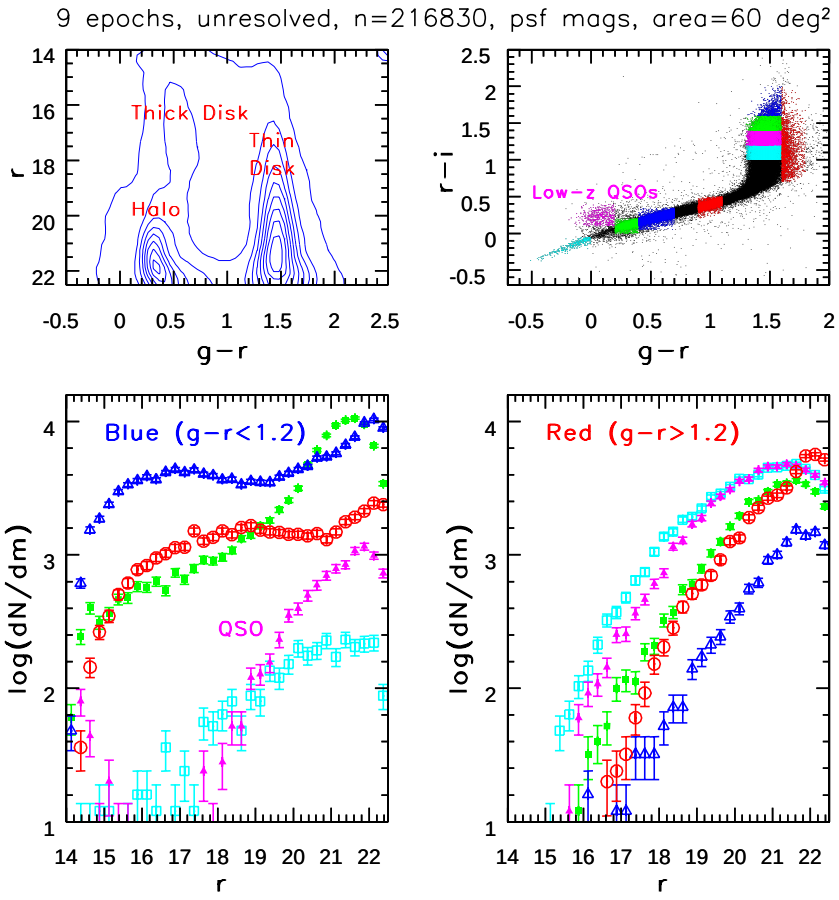

Fig. 8.— Dissecting the Milky Way. The top left shows the distribution of point sources in an SDSS color-magnitude diagram. The main features that reflect the influence of different Galaxy components are marked. The counts of sources selected in the $r-i$ vs. $g-r$ color-color diagram (upper-right panel) are shown in the bottom two panels, showing the bluer (left) and redder (right) stars separately; the color scheme matches that of the color-color diagram. Data such as these as a function of direction on the sky can be a powerful probe of the structure of the Milky Way galaxy.

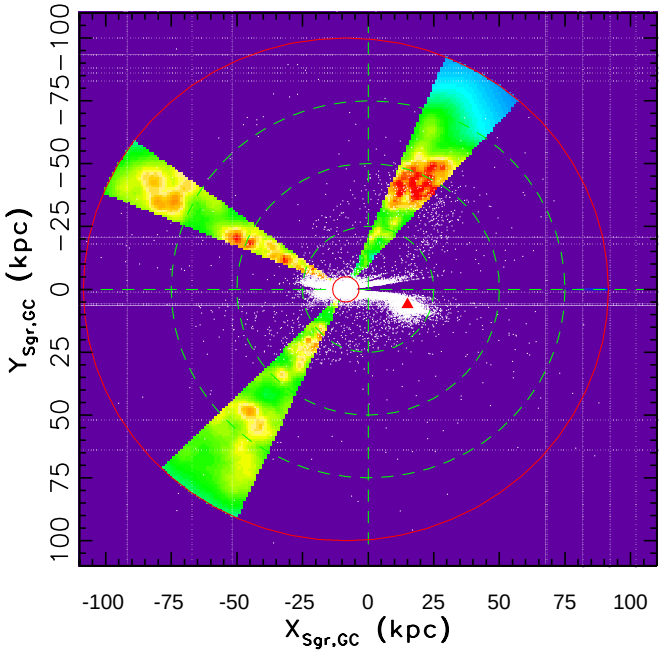

Fig. 9.— The number density multiplied by the cube of the galactocentric radius, and displayed on a logarithmic scale with dynamic range of 1000 (from light blue to red), for 923 SDSS candidate RR Lyrae stars within 10° from the Sgr dwarf tidal stream plane. The solid circles show the sample distance limits (5 kpc and 100 kpc). The dashed circles are centered on (X=0,Y=0), and have radii of 25, 50, and 75 kpc. The triangle marks the position of the Sgr dwarf core. The clumps at (X,Y) of $(20, -35)$ and $(-20, 25)$ kpc are definitely associated with the tidal stream, as is discernible from the distribution of 2MASS M giants (Majewski et al. 2003), shown as the white dots. Other clumps, while consistent with being part of the stream, could also be unrelated super-Poissonian fluctuations, such as those suggested by Bullock et al. (2001). LSST will extend such mapping to about 50 times larger volume.

7. Gravitational Lensing, Weak and Strong

7.1. Introduction

In the last decade, astrophysicists have converged on a standard model of cosmology (e.g., Spergel et al. 2003). While this model explains the current data and is supported by a rich theoretical framework, it invokes two mysterious new components: dark matter and dark energy. The Large Synoptic Survey Telescope can probe these unseen components of the universe with great precision using the subtle distortions of distant galaxies by gravitational lensing. LSST offers the opportunity to take the weak lensing techniques that have been developed on small patches of the sky and extend them to a large fraction of the Celestial Sphere, thereby realizing the promise of these mass maps for high-precision cosmology.

Measurements of weak lensing (WL) shear over a sufficient volume can address cosmology through the measurement of the expansion history of the universe and the growth of structure. The theory and methods for extraction of cosmological information from weak lensing surveys are advancing rapidly at present, but it is already clear that there is tremendous potential for measuring dark energy and dark matter properties. Combining a multiband hemisphere-scale WL survey with the expected *Planck* CMB data could constrain both the dark energy equation of state *and* its time derivative 60 or more times more accurately than the Planck data alone, and measure the neutrino mass to ± 0.05 eV (Song & Knox 2003). These constraints are even tighter (and complementary to) those expected from SNAP SNIa measurements.

Weak lensing measurements on small scales are limited by the noise from the intrinsic ellipticities of galaxies; one wants deep images (especially in g and r, the best bands for shear measurements) to maximize the number of resolved galaxies per unit area on the sky. On larger scales, one is limited by cosmic variance; we therefore need to maximize the sky coverage of the survey. Realizing the potential of a deep, large-area WL survey will require a substantial reduction in the systematic errors arising from atmospheric and instrumental image distortions. Further research is required before we know the limiting factors of an LSST WL survey, but it is clear that large sky coverage with multiple exposures covering each patch will average down many of the statistical errors and allow a variety of tests for systematic errors.

Below we will argue for a baseline LSST WL survey deep enough to obtain meaningful shape measurements of sources out to $z \approx 3$ over of order 15,000 deg^2. Such a dataset, in enough bands to allow photometric redshifts, will enable multiple cosmological tests to be carried out. No other existing or planned observatory would be capable of such a survey. All quantitative statements in this section assume such a survey.

7.2. Weak Lensing and Dark Energy

The physical origin of the recent acceleration of the expansion of the universe is a great mystery and opportunity. Is it "dark energy" or new gravitational physics? In either case LSST is uniquely capable of addressing the underlying physics via all-sky weak gravitational lensing. See also § 8, where the use of supernovae as a probe of dark energy is discussed.

The theoretical community is very actively exploring different cosmological tests using weak lensing, and the list of tests described here will certainly be dated by the time the LSST comes on-line. WL signals have depth information and non-linear, non-Gaussian components, so the full exploitation of WL surveys requires more sophisticated analyses than those already developed for the CMB. But already we know that weak lensing is a tremendously powerful probe of cosmological models, one that is beautifully complementary to those that have given rise to the current standard model of cosmology. Four distinct and largely independent methods of constraining cosmology with WL surveys have already been proposed:

- The number density of clusters as a function of redshift (*"cluster tomography"*). Individual massive clusters provide shear signals large enough to be easily detectable in LSST mass maps. The number of these clusters depends very sensitively on the amplitude of density fluctuations and on the cosmological distance scale, which sets the volume-redshift relation. Both of these inputs depend on the details of dark energy. Hence, the counting of clusters provides a superb way to constrain the equation of state of dark energy.

- The large-angle cosmic shear power spectrum (*"power spectrum tomography"*). This is a quantity that can be measured as a function of cosmic time using photometric redshifts. Combining these results with those from cosmic microwave background data allows constraints on the physics of dark energy. Adding cross-correlations with foreground galaxies increases the precision.

- The non-Gaussian information in the shear field is starting to be explored (*"bispectrum tomography"*). Takada & Jain (2003) show that the bispectrum information from an LSST-scale survey is as powerful as the power spectrum in constraining dark energy parameters. The investigation of such non-Gaussian signatures is in its infancy. Large data sets, such as LSST could produce, are required.

- Shear from foreground sources as a function of redshift (*"shear cosmography"*). The amount of WL shear caused by a foreground structure of known redshift depends on the distance of the background galaxy. Measurements of the shear signal of known foreground structures as a function of the redshift of the background sources (as determined from photometric redshifts) thus allows a measurement of the redshift-distance relation, which depends on the cosmological model. Because this method requires no theoretical estimates of the foreground mass structures, it works even on small angular scales, where predictions of the growth of structure become difficult.

The next few years should see an expansion of the list of cosmological probes using WL. The full utility of the mass maps, in combination with precision CMB data, is not yet understood theoretically, but a clearer understanding of their combined power is emerging.

7.2.1. Cluster Counting via Weak Lens Tomography

The number density of massive galaxy clusters depends very sensitively on the amplitude of density fluctuations and on the cosmic distance scale. Both depend on the details of dark energy. The counting of clusters versus redshift thus provides a superb way to constrain the equation of state of dark energy. The weak lensing shear maps that LSST will provide can be turned into maps of projected mass density. Using shear data for source galaxies at different redshifts allows determination of the detected cluster redshift (see Figure 10). Weak lensing measures the mass of the cluster directly, thereby removing one of the main uncertainties in using clusters of galaxies as a cosmological probe. This is particularly important, as the number density of clusters is an exponential function of mass. Because of this sensitivity, accurate, *direct* measurements of mass such as weak lensing provides is particularly important.

With a sky coverage of 15,000 deg^2, the LSST is expected to find about 200,000 clusters in its mass maps; a cluster sample of this size offers 2-3% precision on the equation of state parameter w if a degeneracy between w and Ω_M is broken using other probes (Tyson et al. 2003).

The shape of the observed distribution of clusters with mass and redshift needs to be compared with N-body simulations to place quantitative constraints on cosmological models (e.g., Hennawi et al. 2004). Cluster counting achieves its best accuracy at $z < 1$, where the effect of dark energy on the volume-redshift relation is maximized.

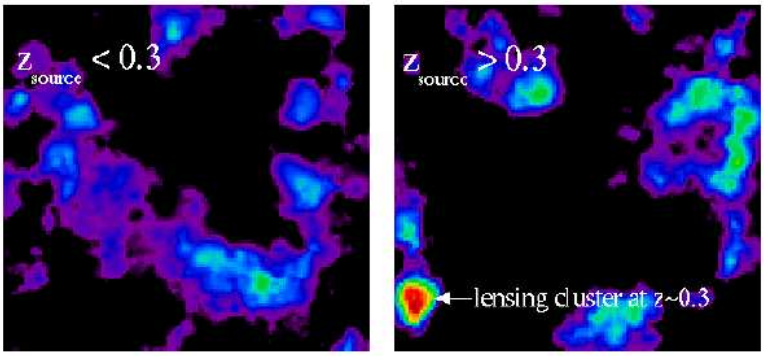

Fig. 10.— 3-D cluster tomography from the Deep Lens Survey (Wittman et al. 2001,2002). These mass maps of a $40'$ field show two slices in redshift. Similar 3-D mass tomography has found clusters up to $z = 1$.

Sensitivity to the growth function and volume element make this method complementary to "metric" methods such as supernovae (Riess et al. 2003) or redshift surveys for acoustic peaks (e.g., Seo & Eisenstein 2003).

Cluster counting is relatively immune to systematic errors in shear measurement. Additive errors, e.g. from uncorrected PSF variation, act as a small additional noise source since they will not be correlated with the true cluster distribution on the sky. Calibration errors, e.g. photometric-redshift biases or uncorrected "circularization" of galaxies by the PSF, will produce small mass-scale miscalibrations. Projection effects can bias cluster masses upwards, but this can be quantified with aid of N-body simulations.

7.2.2. Power Spectrum and Bispectrum Tomography

Using the same data as a WL cluster survey, an independent probe of dark energy comes from weak lens shear tomography, in which shear of sources in various redshift bins is correlated over wide angles. The expected power spectrum of WL is given by (Refregier 2003):

$$C_l = \frac{9}{16}\left(\frac{H_0}{c}\right)^4 \Omega_m^2 \int_0^{\chi_A} d\chi \left[\frac{g(\chi)}{a r(\chi)}\right]^2 P\left(\frac{l}{r}, \chi\right) \tag{2}$$

where P is the power spectrum at comoving distance χ, r is the angular diameter distance divided by $a \equiv 1+z$, and $g(\chi) \equiv 2\int_\chi^{\chi_A} d\chi' n(\chi') \frac{r(\chi)\, r(\chi'-\chi)}{r(\chi')}$ is a radial weight function for lensing for a sample of galaxies with density n as function of distance. Thus it depends both on the growth of matter fluctuations and the angular diameter distance relation; both of these terms are dependent upon the quantity and equation of state of the dark energy, and the former is sensitive to the dynamical properties of the dark matter. Using photometric redshifts to characterize the lensing signal as a function of source galaxy redshift, one can measure the growth and geometry as a function of redshift. Doing so gives much more sensitivity than using the full 2-D projected power spectrum (Hu & Keeton 2002).

Figure 11 shows an estimate of the precision achieved in cosmological parameters in a cosmic shear survey of 15,000 deg^2 to sources out to $z = 3$. This is taken from the work of Hu & Jain (2003). A density of 70 sources per arcmin^2, a floor to the shear errors of 0.0001 (cf., § 7.4.3), and the WMAP CMB priors were used. The effect of including two-point galaxy correlations is also shown.

Figure 12 shows the large-scale two-dimensional power spectrum at various redshifts, together with the expected errors for a tomographic survey covering half the sky. At large scales (l less than a few hundred), where the power spectrum and bispectrum are most reliably predicted with perturbation theory, the measurement is limited

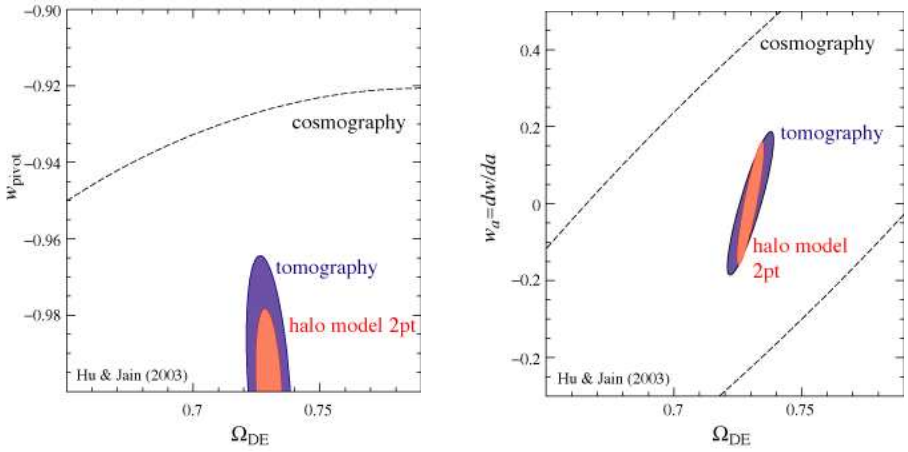

Fig. 11.— One-sigma error ellipsoids in the planes of some dark energy parameters for a 15,000 deg^2 weak lensing survey; see Hu & Jain (2003) and the text for more details. Left: dark energy equation of state w at $z_{pivot} \approx 1/3$, versus the present-day contribution of dark energy. Right: Limits on $w_a \equiv dw/da$ and dark energy fraction. These constraints do not make use of cluster counting or bispectrum information.

by sample variance. Also shown is the effect of a floor in the systematics of the shear due to uncorrected PSF variation at the level of 0.0001; it is smaller than the expected signal at all l.

The non-Gaussian bispectrum signal is likewise limited by sample variance at $l \lesssim 1000$, hence a full-hemisphere survey is required for useful constraint of the large-scale information. LSST can measure the cosmic shear power spectrum with precision, addressing the physics of dark energy.

The universe may have additional surprises for us regarding the evolution of cosmic perturbations. For example, non-zero neutrino masses or other small admixtures of hot dark matter would cause the LSST-measured shear power spectrum to be offset from the CMB prediction. The WMAP data show a marginally significant deficit of power at the largest measurable scales ($l < 10$); these measurements will not improve, because we have surveyed the entire last scattering surface. With very-large-scale WL measurements one can, however, measure the power on similar comoving length scales over the *interior volume* of the last scattering surface, and potentially reduce the cosmic variance enough to conclusively test for the deficit of large-scale power.

For cosmological tests involving anomalies in broad-band shape and amplitude of the power spectrum, weak lensing measurements will become the method of choice in the coming decade. Only by probing the dark matter directly with weak lensing will we be able to detect physical effects that induce one-percent-level changes in the power spectrum.

Power spectrum and bispectrum measurements are sensitive to additive systematic errors from uncorrected PSF distortions, because the power from the systematic error will add to the real power. Shear and photo-z calibration errors at the $\sim 1\%$ level will also become important because the S/N level on the bandpower estimates will be $\gtrsim 100$ on all but the largest angular scales.

7.2.3. Shear Cosmography

The shear signal from gravitational lensing depends on the relative separations of the source, lens, and observer. By comparing the differential signal from foreground structures of known redshift between background sources at multiple redshifts, one can achieve precise measurements of the distance-redshift relation (Jain & Taylor 2004;

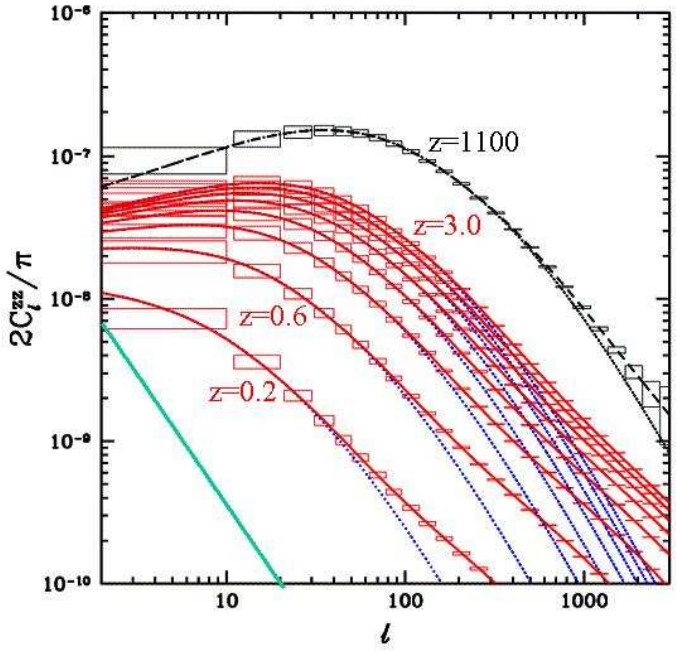

Fig. 12.— The solid red curves show the shear power spectra of galaxy sources divided into eight different source redshift bins, starting with the bin centered at $z = 0.2$ with the least power, rising to $z = 3$ with the most power. The error boxes are the forecasted level of error assuming coverage of half the sky with a total source density of 65/sq. arcmin. The black curve is the shear power spectrum for a source at $z = 1100$, the CMB last-scattering surface. CMB error boxes are forecasts for a shear power reconstruction from a low noise, high resolution all-sky CMB polarization map. Dotted curves are predictions of linear theory, solid curves include non-linear effects. The green curve shows the contribution from an uncorrected shear systematic of 0.0001. Figure courtesy of Lloyd Knox.

Bernstein & Jain 2003), yielding a measure of the dark energy equation of state. This method differs from the shear power spectrum in that it does not attempt to characterize the statistics of the mass distribution of lenses, but focuses merely on measuring relative cosmological distances. Because there is no attempt to characterize the statistics of the mass distribution, there is no sample-variance limitation to this method, and the cosmological sensitivity is controlled primarily by the total number of usable source galaxies in the survey. The lensing signal can be used even at very small scales where N-body simulation of the growth of structure is impractical.

For an LSST-scale survey free of systematic errors, the dark energy constraints from cosmography are comparable to those expected from the SNAP SNIa survey. More importantly, they have different degeneracies from the SNIa, power spectrum, CMB, and other measures, and hence they substantially improve the accuracy of the combined ensemble of measurements. In addition, we've seen that weak lensing gives a variety of *different* techniques to constrain the nature of dark energy, each with different systematics, and thus providing a series of powerful cross-checks on the results.

The cross-correlation between foreground structures and background shear makes this method quite insensitive to systematic contamination of the WL shape measurements, as they are uncorrelated with the foreground mass. On the other hand, the method is very sensitive to small redshift-dependent errors in shear and photo-z calibration, because one is looking for small changes in the ratio of shears at different redshifts. Hence the calibration demands are more stringent than for cluster-counting or power-spectrum tomography.

7.2.4. Accuracy of Derived Parameters

Combining data on the anisotropy of the CMB from WMAP and Planck with weak lensing data from LSST should yield tighter constraints on the nature of dark energy. Combining the two measurements will pin down w to of order 3% and its time derivative w_a to 0.05, as well as tightly constrain the dark matter power spectrum (Hu & Jain 2003). These constraints combine those from power-spectrum tomography, CMB data, and limited use of the cosmography method—they do not include possible information from cluster counting or non-Gaussian statistics. Inclusion of SN Ia information would further tighten the constraints. These errors are for a simultaneous fit to a full set of cosmological parameters, and include marginalization over all parameters except curvature.

Even if the CMB constraints are not included, Ω_m and other parameters are tightly constrained. This overdetermination allows for consistency checks between WL, CMB, and SNIa methods.

7.3. Gravitational Lensing and the Nature of Dark Matter

The physical nature of dark matter is unknown, although some candidates have been ruled out via astrophysical data. Aside from possible future detection of WIMPS or axions in the laboratory, the strongest constraints on the physics of dark matter come from the way it clumps on small scales. For example, Monte-Carlo simulations of the clumping of Cold Dark Matter are at variance with current observations of the distribution of mass on galaxy to cluster scales. High resolution N-body CDM simulations (e.g., Moore et al. 2000) show steeper inner mass profiles than observed (Tyson et al. 1998; Sand et al. 2003). This has led to speculation that dark matter may be self-interacting (Spergel & Steinhardt 2000, Firmani et al. 2000).

The most sensitive and direct way to map dark matter on small scales (1-100 kpc) is via strong gravitational lensing. To explore cosmic variance and astrophysical effects, many examples of strong lensing from galaxy to group to cluster scales are required. Such systems are rare, but the deep multi-color survey covering much of the visible sky that LSST will perform will find many examples.

7.3.1. Dark Matter in Cluster Cores

Of the 2×10^5 or so galaxy clusters discovered and mapped in the weak lensing survey, of order 1000 should exhibit luminous arc systems. The strong lensing constraints will greatly enhance the precision of the mass estimates of this cluster subsample, and allow the dark matter halo profile to be measured over a larger range of length scales. Figure 13 gives an impression of the image quality to be expected from the LSST cluster images. Colley et al. (1996) obtained a high-resolution map of the dark matter in this cluster, which was possible because one of the multiple images passed near the center of the cluster lens; such alignments are rare, and we expect only 50-100 such systems in the whole sky. N-body CDM simulations of structure formation suggest that the inner mass profile may show large scatter (Fukushige et al. 2003); in order to explore this, we need a large sample of clusters. All of this argues for a very wide, deep survey for strong lenses.

7.3.2. Galaxy-scale Dark Matter

The CLASS survey (Browne et al. 2003) has found that 10^{-3} of extragalactic radio sources, i.e., random lines of sight, are lensed. With 60 sources per arcmin^2, we therefore expect of order $\approx 10^6$ multiple-image systems—mostly galaxy-galaxy lenses—to be present in a 15,000 deg^2 survey. Recognizing these objects as lenses will require more work, however. The lensing cross-section is dominated by massive elliptical galaxies at redshifts $0.3 < z < 1$

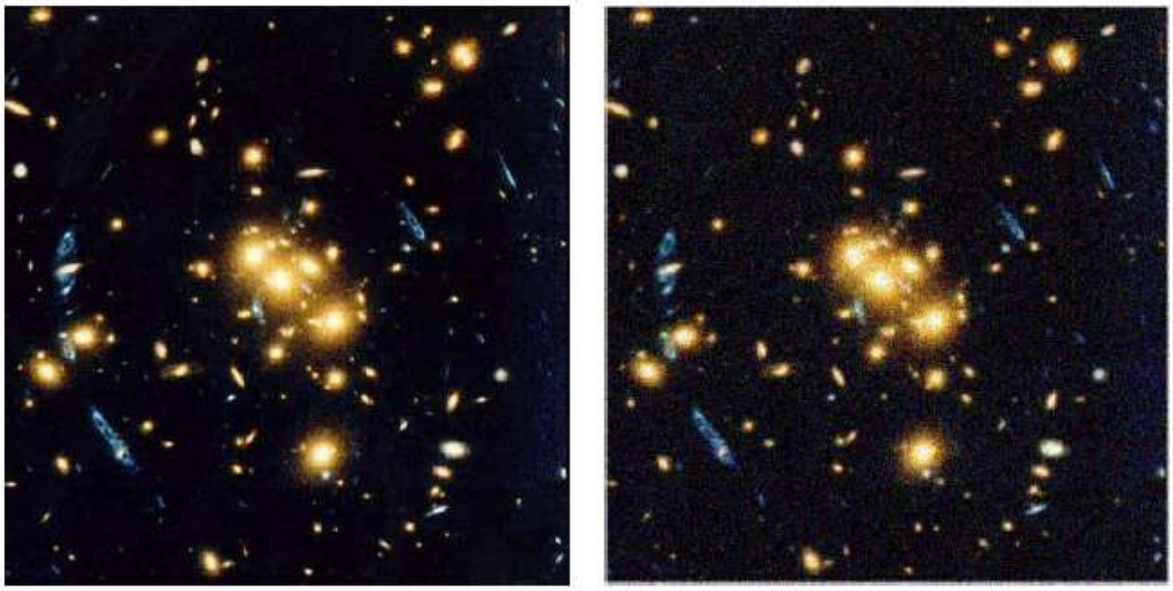

Fig. 13.— Left: HST image of the inner regions of CL0024+1652 (from Colley et al. 1996). Right: A simulation of the LSST view of the same field, generated by convolution with a seeing disk of width 0.7 arcsec, re-binning into LSST pixels, and degrading with a small amount of noise. The lensed images are still partially resolved, providing information about the lens structure on small scales.

(e.g. Fukugita & Turner 1991); scaling from the Hubble Deep Field, we may expect approximately 10,000 such "clean lens" galaxies per deg^2; if each has an Einstein ring radius of about $1''$, this adds up to ~ 30 deg^2 of lensing cross-section in the whole survey. These galaxies can be used as markers for strong lenses; one can, for example, search for achromatic excesses in their images. Note the great importance of multi-color LSST imaging in this task.

Large samples of galaxy lenses will permit new tests of the structure of galaxy dark halos. For example, it is believed that small-scale substructure in halos leads to anomalous flux ratios in multiple-image systems. With a large sample of four-image lenses, these flux ratio anomalies, and hence the halo substructure, could be carefully quantified. Within the CDM framework, the abundance of wide-separation ($> 3''$) lenses is a sensitive test of the shapes of CDM halos. The SDSS has discovered one very wide QSO system ($14''$; Inada et al. 2003; Oguri et al. 2003), but an LSST survey would tremendously improve the statistics.

The very high magnifications in rare lenses provide us with a very powerful "cosmic telescope". The LSST optical data on sources observed in this way will provide very important complementary information to that available in similar scale low cadence surveys across the rest of the electromagnetic spectrum, including those by EXIST in the X-ray band and the Square Kilometer Array in the radio.

7.4. Image Quality and Weak Lensing Analyses

The theoretical analyses today show that the weak gravitational lensing pattern on the sky holds vast amounts of information on dark energy, neutrino masses, and other cosmological quantities of interest. The WL signal is, however, quite subtle, amounting to $\approx 1\%$ distortion of the typical background galaxy shape. This is already below the 1–10% asymmetries seen in stellar images from a typical telescope. This "foreground" signal due to the point

spread function (PSF) of the atmosphere, telescope, and detector, must be removed to high accuracy in order for us to detect subtle cosmological signals that are but a small fraction of the total lensing signal. Fortunately, the stellar images in the survey fields themselves serve as test points for optical distortions. We can take the PSF information derived from the stellar images, and use deconvolution algorithms to extract the intrinsic (pre-seeing) moments of the galaxies from their observed shapes.

In present WL surveys, the residual errors from this PSF-correction process are as large as the random errors due to "shape noise." An LSST survey, with $\approx 10^3$ times as many galaxies as present surveys, will have to reduce systematic errors by 3–10 times from presently achieved levels in order to avoid being dominated by PSF systematic errors. This is a significant challenge that can be addressed by improved deconvolution algorithms (Bernstein & Jarvis 2002; Refregier & Bacon 2003), and by designing telescopes and surveys explicitly to meet this challenge. Further study and experimentation are required in order to accurately estimate the systematic-error floor for an LSST WL survey. Here we summarize the most important sources of error and means of detecting and eliminating these errors.

7.4.1. The Effects of seeing

The point spread function (PSF) delivered by the atmosphere and optics most strongly affects weak lens shear measurements. First, there is systematic error from the shape of the PSF on scales uncorrectable by stellar PSF data. This effect primarily additive in the sense that it imposes an artificial shear pattern (that of uncorrected PSF ellipticities) atop the true lensing shear pattern. Second, there is the loss of angular resolution on compact galaxies. This is primarily a multiplicative effect in that the PSF "circularizes" the true galaxy isophotes to some degree that must be corrected by the deconvolution. It also increases the noise level, as the intrinsic shape of a poorly-resolved galaxy cannot be meaningfully recovered from an observation with finite S/N.

Both of these effects may be reduced by going into space, where very small and stable PSFs are available. One can also obtain higher-S/N images of the galaxies, which improves the quality of the deconvolution, although there will always be some minimum galaxy size (relative to the PSF) below which a reliable shape measurement is not feasible. The mean FWHM of galaxies as a function of redshift is larger than 0.6 arcsec. A large total aperture increases the sensitivity to low surface brightness, a requirement for resolving and measuring the shape of high redshift galaxies.

The first effect, spurious PSF power, can be mitigated in part by stacking hundreds of exposures, thus averaging down the unmodeled variations in PSF shape. This meshes very well with the expected LSST survey strategy.

7.4.2. Targets for Systematic Error Levels

The desirable levels of additive and multiplicative systematic errors in the shear maps can be estimated from Figure 12; the plotted quantity $2C_l^{zz}/\pi = l(l+1)C_l/2\pi$ is the lensing shear variance contributed per log interval in angular scale. The sample variance in C_l in a logarithmic band of width l is roughly $(\Delta C_l)^2 = \frac{C_l^2}{l^2 f_{sky}}$, where f_{sky} is the fraction of the sky covered by the survey. At degree scales ($l \approx 10^{2.5}$), $l^2 C_l/2\pi \approx 10^{-4.5}$, giving an RMS shear somewhat below 1%, thus the cosmic error in C_l is of order 10^{-7}. We wish that systematic errors in the shear measurement (due to, e.g., uncertainties in the PSF) be below this, i.e., well below ~ 0.0002. Likewise any shear calibration errors should be held below about 0.5%. At larger angular scales, l =10–100, the expected WL bandpower is lower, but the sample variance is higher, leading to only slightly more stringent systematic-error variance limits on larger scales. The calibration demands are relaxed on these larger scales.

Present-day surveys offer some guidance to the required advances in technique. All surveys to date have systematic-error band-powers that are $\sim 10\%$ of the cosmic WL signal at the most favorable angular scales. At less favorable angular scales, the systematic-error signals are as large as the WL signal. Thus the PSF-residual variance is $\approx 10^{-6}$ in surveys of brighter (and hence well-resolved) galaxies, suggesting that a factor of 3–10 reduction in RMS residual PSF signal would suffice. Significant further development and simulation are required in order to determine how well one can do for fainter (hence more poorly resolved from the ground) galaxies and on very large angular scales.

The quality of LSST seeing is a prime determinant of the limits of the WL survey. Better seeing means that more galaxies could be resolved sufficiently well to use for high-accuracy lensing surveys. Data from the Deep Lens Survey (e.g., Wittman et al. 2003) go to $R \sim 26$ in 0.7–1.0″ seeing. At 1″ seeing, with a depth of $R = 25.5$ for 5σ detection of point sources, 35 galaxies per arcmin^2 are resolved enough to be usable for shear measurements. Better seeing will undoubtedly increase the number of useful galaxies (and hence reduce the random shape noise). These numbers may go as high as 60 galaxies per arcmin^2 in 0.7″ seeing. What is the expected number of sources per square arcminute and the effect of seeing in the ground-based LSST?

We can investigate the effect of seeing on tangential shear from a typical cluster of galaxies (700 km/sec velocity dispersion) at $z = 0.3$. A model cluster with a soft core and an NFW outer profile was placed in front of a mosaic of HDF deep images assumed to be at $z = 2$ (24-27 mag galaxies were used). No PSF corrections were made. A range of seeing was simulated, with PSF FWHM from 0.3 to 0.7 arcsec. A gradual degradation of the measured shear is seen, correctable via scalar PSF shear "polarizability" correction, with no evidence of systematic effects which could affect the inferred distribution of clusters in mass and redshift. Over 90 galaxies per square arcminute were found in 0.7 arcsec seeing, while 70 per square arcminute survived the pipeline filter for overlap projections, and 60 are resolved enough to allow shape measurements to be performed. This simulation compares favorably with recent deep data from imaging on Magellan: For a 60-minute stack of R-band images taken with the MagIC direct imager on Magellan II during mean seeing of 0.7", 85 galaxies per arcmin^2 are photometered down to r 26.

7.4.3. Controlling shear systematic error

The LSST can be designed to significantly assist the reduction of WL systematics from PSF residuals. The best current WL surveys have residual PSF systematics at roughly the 0.001 level of shear in the coadded image stack. This is adequate for cluster detection but not for low level cosmic shear correlations of the sort envisioned for LSST. While atmospheric seeing and differential chromatic refraction will always be (time-variable) contributors to the PSF, existing telescopes were not designed to minimize aberrations or time variations in telescope PSF due to mirror motion under gravity (10% PSF ellipticity change between exposures is not uncommon). Faint stars in the field are used to map the delivered PSF variations in each exposure, but the low density of stars in high-latitude fields means that PSF variations cannot be tracked if they vary rapidly in space and time. *Thus the key is control of PSF shape changes between exposures on sub-arcminute scales.* The LSST telescope and camera, and the observing strategy, must be designed from the ground up to minimize many of the problems which plague current facilities. The science goals of § 7.2 require a systematics floor "only" a factor of five lower than that obtainable with the old technology 4-meter telescopes built thirty years ago. Further study is required to determine the ultimate limits of PSF stability from the atmosphere and a well-engineered telescope. By chopping the shear signal in multiple ways we will beat down the systematics even further. The goal of a systematics floor of 0.0002 shear should be attainable for LSST.

7.4.4. Diagnosing Systematic-Error Contamination

After we have constructed a telescope of ideal design and used the best possible deconvolution algorithms, there are several ways in which we can test for the presence of uncorrected systematic signals in the measured shear maps. First, the shear field is a two-component observable that is derived from a scalar field, the projected gravitational potential. It is therefore possible to decompose the shear field into two components, called "E" and "B" modes, with all the gravitational signal confined to the E mode. The B mode is a diagnostic for uncorrected systematic errors, which will not in general be confined to the E mode.

Additional cross-checks for systematic errors would be enabled by the ability to view the same field at different rotation angles relative to the detector or telescope axes, either because of the natural field rotation on an alt-az telescope mount, or through an ability to rotate the camera assembly on the telescope.

A further important point is that the WL analysis methods described above—cluster counting, power spectrum, bispectrum, and cosmography—are each sensitive to systematic errors in different ways. The cosmography method, for example, is nearly insensitive to additive errors, but highly sensitive to multiplicative or photometric-redshift errors, compared to power spectrum or bispectrum methods. Multiple analyses within the LSST data set can provide cross-checks on cosmological inferences. Space-based weak lensing data would have very different characteristics as well—lower systematic errors, but larger random errors—enabling further cross-checks. These weak lens techniques are subject to very different systematics from other measures, such as CMB, SN Ia, or "baryon wiggle" methods, and their degeneracies in the cosmological-parameter space are quite distinct as well. The ensemble of measurements should be quite robust.

7.4.5. Photometric Redshift Errors

The tomographic techniques require estimation of photometric redshifts for all galaxies in the survey—both source and lens galaxies. Photometric redshifts allow us to understand the redshift of the cluster lenses as well as the differential shear of multiple source redshifts. Fortunately, the photometric redshifts need not be superbly precise, even 5–10% in $(1 + z)$ would do, but extending the range of redshifts over which the results are robust is very important. Furthermore it is important that the *bias* in photo-z data be much smaller than the errors on individual galaxies, so that the estimates of the ensemble redshift distribution are highly accurate. This will require well-selected filter bands and a substantial, but feasible, parallel program of spectroscopic redshifts to act as a calibration set. Five bands will be required. Finally, while the required Gaussian dispersion in the photometric redshift relation of 5-10% in $(1+z)$ is achievable from current ground-based surveys, the requirements on the controlling the non-Gaussian tails of this dispersion are an order of magnitude more stringent than surveys currently achieve, and this will be a challenge.

7.5. Turning Science Goals into Requirements

We have outlined the science potential and leading error sources for the gravitational lensing measurements that would be possible with an LSST. Translating science goals into engineering requirements requires further detailed examination of systematic errors in the WL data as a function of PSF size and stability, tradeoffs in depth vs area, and the filter combinations that optimize shape and photometric-redshift measurements. In this section we give some preliminary estimates of the specifications that will be needed to extract the scientific potential of ground-based WL surveys.

7.5.1. Shear Measurement

Weak lensing requires both a deep survey and one covering as much of the sky as possible. Large-scale power-spectrum and bispectrum measurements are limited by sample variance, which pushes the survey to cover the largest possible area of sky. Smaller-scale spectra, and "cosmography" methods, improve as the total number of galaxies with measurable spectra. For an instrument with the throughput of LSST, this drives one to spread the available integration time over essentially the full available high-latitude sky.

Cluster counting constraints scale as the volume surveyed, which again drives toward a full-hemisphere survey. N-body simulations suggest that one should measure $> 200,000$ massive clusters distributed out to $z = 1$, which would require a $\approx 15,000$ deg^2 survey.

For a ground-based survey, the atmospheric seeing limits the number of resolvable galaxies per square arcminute to of order 60, as we've seen above. Once these galaxies are measured at $S/N \gtrsim 20$, there is little additional shape information to be gained, hence there is a natural limit to the desirable depth of ground-based WL measurements. This limit is likely near $R \lesssim 26.5$ mag.

The PSF quadrupole moment must be stable on arcminute scales at the 1% level during an individual exposure, so that it can be corrected from the measured PSF of (relatively sparsely sampled) stars. For control of systematic PSF errors, repeated short sky-limited exposures are necessary. Having the hundreds of exposures of each field over a variety of position angles will further reduce systematic effects.

7.5.2. Photometric Redshifts

The primary leverage for addressing dark energy is in the redshift dependence of the shear field. Distance information for tomography comes primarily from photometric redshifts of the source galaxies. This requires multi-color imaging in at least five bands stretching over as wide a wavelength range as possible. Going as red as the Y band would be helpful as it would extend the reach of diagnostics such as the 4000Å and Balmer breaks out to $z \sim 1.3$, and is more effective than a U filter for high-redshift galaxies. Photometric redshifts of 1σ accuracy $0.1(1+z)$ are adequate for the science goals described in this section, and require a 10σ photometric accuracy in 5 bands for galaxies brighter than $r_{AB} = 26.5$. For comparison, current ground-based 4-band data at $r = 25.5$ AB mag achieves 6% precision in $1+z$ out to $z = 1.3$.

This work is critically dependent on the accuracy of photometric calibration; the calibration must be uniform to 2% rms (1% is probably feasible, and should be the aim). Precise requirements on the calibration remain to be determined from photo-z modeling studies. Photometric redshift accuracy will also be impacted by systematic photometric errors due to problems in deblending, cosmic rays, and so on; one must include these effects in the error budget. Note that weak lensing makes fewer demands on the *image quality* of the bands other than the principal bands in which the shape measurements are done, as long as one achieves the necessary depth. In addition, weak lensing makes no demands on the cadence of observations; observations at all epochs can be added together. Indeed, adding observations made at different times of the year and different camera rotations helps to reduce systematic errors in determining galaxy shapes. Short exposures optimize reconstruction of the high resolution image stack.

7.5.3. Cluster dN/dz

Measurements of the redshift dependence of the cluster mass function will require roughly 5 z-bins, 10 mass bins, and at least 10 fields on the sky each with thousands of clusters (for cosmic variance). Based on N-body

cosmological simulations, this would require about 10,000 clusters in the most massive bins for each field: over 200,000 mass clusters distributed up to $z = 1$. This requires at least 15,000 deg^2 in a shear survey down to about 26th r mag equivalent in at least 5 optical bands, the same requirements we've found above. While the shear generated by clusters of over 10^{14} solar masses is large compared with the typical systematic shear errors, one must be sure that there is no bias versus redshift of $M_{min}(z)$ propagating to systematic errors in N(M,z). One worries about such a bias because the minimum detectable mass at a given S/N ratio is a function of redshift.

The mass function (for modern CDM models and their variants) is universal only if we measure mass within a rather large radius of about 3 Mpc. At this distance from a 10^{14} $M_\odot$ cluster at $z \approx 1$, the shear is 0.001; we will want systematics to be substantially smaller than this. Of course, the shear will be larger on smaller scales, where shear error from PSF systematics is more difficult to control. Thus, the shear systematics required for cluster counting are less stringent than for cosmic shear. The main requirement is a sufficiently large sample of clusters spanning the mass function and redshift range, so that the sample can be cut in multiple ways.

7.5.4. General instrument specifications

All of the above leads to a series of requirements on the telescope, instrument, and observing strategy:

- *Flux limits and survey area required:* The area of sky imaged at any given time must be as large as possible consistent with the PSF specifications below. This maximizes the number of overlaps a given area of sky will receive, which is crucial for control of shear systematics to the 0.0002 level.
  - *The total area of sky to be covered:* Roughly 15,000 deg^2.
  - *The depth and dynamic range needed in a single exposure:* Sufficient to cover photometric and astrometric transfer standards (roughly 17 mag).
  - *The depth and dynamic range needed in stacked exposure:* $r > 26.5, B = 26.5, g = 26.5, i = 26, z = 25$ AB mag ($10\,\sigma$) over survey lifetime.
  - *Length of individual exposures:* This is set by the dynamic range requirement above, the properties of the detectors, and the telescope throughput. Roughly 15 seconds for the 8.4m LSST design.
  - *Requirements on slew time:* Slew time must be kept short in order to maximize observing efficiency; approximately 5 seconds.
- *Requirements on PSF and pixel size:* The FWHM of the PSF should be less than roughly $0.7''$ median in the bands in which shear measurements are made (r, g, and/or i). Time and spatial variations of PSF between exposures, particularly its ellipticity, must be smaller than 0.01 on scales of the mean distance between faint stars (20–40$''$). After correction from the measured PSF of stars, the systematics in the PSF ellipticity should be at the 0.001 level or less. Other bands used for photometric redshifts benefit from good FWHM but have no specific PSF ellipticity specifications. In order to properly sample the PSF for stars, especially in the best seeing, we require pixels of $\approx 0.2''$.
- *Filters and cadence required:* For shear measurements, mostly r, g, and i images, returning to the same patch again on the timescales of changes in optics and atmosphere PSF systematics. Photometry in *bgriz* and possibly Y (time permitting) for photometric redshifts. The Y filter is centered at 1 micron, beyond where CCDs have traditionally been sensitive, but modern devices have sensitivity curves that only fall drastically at 1.1 micron. The longer-wavelength bands (z and Y, and perhaps i) could be carried out in moonlight. The cadence described in Appendix A is adequate for this job. Initial results from an LSST

observations simulator suggest that an LSST of étendue 270 m^2 deg^2 is capable of covering 17,000 deg^2 to the required depth in *bgriz* over ten years, under realistic models for weather and slew efficiency.

- *Photometric calibration accuracy required:* Photometric calibration of 0.02 mag rms in all bands used for photometric redshifts. This needs to be further quantified, in terms of limits on angular structure in photometric zeropoint errors; however, an observing strategy that involved massive redundancy and overlap should minimize such structure. The astronomical community's experience with exquisitely calibrated photometric data is quite limited. 0.01 mag calibration may be reachable, and we may find ourselves deciding that this level of accuracy is needed.

- *Astrometric accuracy required:* Absolute astrometric accuracy is set only by the need to capture astrometric secondary standard stars (roughly 0.1 arcsec). The relative accuracy is set by the need to measure the field distortion variations between exposures and must be consistent with an induced shear less than 0.0002 on the smallest useful scale. For $l = 2000(360'')$ this corresponds to an error of 80 mas rms per coordinate on each exposure.

- *Tails of the astrometric and photometric error distribution:* Data frames with astrometric errors in the tails can be clipped. Tails of the photometric error are more problematic. Catastrophic photometric redshift errors, if not rejected, bias the mean redshift of a source redshift bin. Systematic effects can arise in faint galaxy photometry due to surface brightness dimming, and this needs to be guarded against.

- *Sky darkness and photometricity required:* b and g observations need to be carried out with the moon below in the horizon. r and i observations can be done with a little bit of moon, and z and Y observations can be done with quite a bit of moon (although this needs to be properly quantified). Photometric conditions are required at least once per season per sky patch. If each patch is imaged 3 times per month, that leads to $\approx 10\%$ minimum photometric nights. Sky brightness in dark time must be similar to that of the best known sites.

- *Auxiliary data required:* A program of spectroscopic calibration of the photometric redshift relation will entail up to a few $\times 10^4$ spectra of 25 mag galaxies. This is at the limit of capabilities of projects such as DEEP2 on Keck, or surveys with VIRMOS on the VLT. We will also need a network of photometric transfer standards, which will be a huge effort, especially if we're using a filter such as Y, which has not been widely used in the past.

- *Specialized data analysis tools needed to carry out the science:* Considerable development of the algorithms for shear extraction, E/B mode extraction, and image co-addition have occurred over the past decade. However, the demands of the LSST WL science are such that further development will be required.

7.6. Where will WL cosmology be in ten years?

Several WL surveys will have been completed in this decade. The Deep Lens Survey (DLS; cf., Wittman et al. 2003), a four band photometric redshift survey to 26th magnitude in six fields covering 24 deg^2, should be completed in 2004 and the analysis and interpretation completed by 2006. This will yield shear correlations out to ~ 1 degree and counts of mass clusters out to $z \sim 0.8$. The Megacam CFHT Legacy Survey has begun and should be complete and largely analyzed by the end of the decade. The WL part of CFHTLS will cover up to 170 deg^2 with photometric redshifts (*ugriz*) in several $6^\circ \times 6^\circ$ fields with depth comparable to the DLS. These surveys will break the $\Omega_M - \sigma_8$ degeneracy independent of other probes, will pin down our location (in current cosmological

models) in the $\Omega_{DE} - \Omega_M$ plane to perhaps 20%, and will catalog about 2000 massive clusters out to $z \sim 0.8$. However, the total number of sources will be a factor of 10–100 smaller than that required to unambiguously measure time evolution of dark energy. Larger surveys will be required to adequately address the physics of dark energy and neutrino masses.

The VISTA (*"Visible and Infrared Survey Telescope for Astronomy"* 4-m telescope will initially be equipped with a wide-field (one square degree) near-IR camera (with the hope of building an optical camera in the future), which will be used for a wide range of surveys; weak lensing studies was identified early on as one of the main drivers of the project. Given the substantially higher sky background in the near-IR, the flux limit at which shape measurements are possible for a given amount of integration time will be substantially higher than for LSST. The science goals of VISTA do overlap quite a bit with LSST, but it is unlikely to be seriously competitive. At least 75% of the time is to be used for large scale surveys. They envisage deep, medium, and wide surveys (probably carried out in time of poorest seeing) each addressing different science goals. Survey operations are expected to begin in 2007. It would be useful if LSST and VISTA coordinated their operations, to optimize the joint multi-wavelength science.

There likely will be smaller (∼degree) areas surveyed more deeply, such as a Subaru survey and an HST ACS survey. If complemented by near-infrared ultra-deep photometry, such pencil beam surveys will usefully constrain the time development of dark matter structure at $z > 1$. Part of the contribution to the cluster counts in the wide surveys, particularly at $z > 1$, comes from this growth function, and these deeper surveys will serve to calibrate the growth contribution to $N(M, z)$. LSST's unique potential contributions to WL cosmology stem from (1) deep multiband coverage of over 15,000 deg^2 yielding the needed statistics to support percent level precision on w, and (2) attention to shear systematics built into the observatory design.

7.7. Comparison of Facilities

We have seen that the WL science goals favor a survey with photometric redshifts and shape measurements for galaxies down to $R \approx 26$ over 15,000 deg^2 of sky.

Like LSST, the Supernova Acceleration Probe (SNAP), a candidate for the NASA/DoE Joint Dark Energy Mission (JDEM), is being designed for optimal performance in weak lensing measurements. The space platform offers a PSF ten times smaller than the ground survey, and exceptional stability of the PSF due to the (nearly) complete absence of atmospheric variation, gravitational loading, or thermal variation in the observatory. Hence one can expect SNAP to usefully resolve and deconvolve at least twice as many galaxies per square arcminute, with lower PSF systematic residuals. SNAP will furthermore obtain simultaneous data from 0.4–1.6 microns, for excellent photometric redshift determinations.

The throughput of SNAP would be $\sim 200\times$ larger than that of the HST Advanced Camera for Surveys (ACS). The LSST, however, will be able to survey far more sky than SNAP: the baseline scenario for SNAP envisions a one-year weak lensing survey of 1000 deg^2, with larger surveys possible in an extended mission. For sample-variance-limited pursuits such as cluster-counting and large-scale power shear power spectra, the proposed LSST would offer $\approx 4\times$ lower parameter errors. For galaxy-count-limited pursuits, such as small-scale power spectra and cosmography, the LSST hemisphere-scale survey has 2–3 times more galaxies. [In fact the parameter limits from SNAP are a bit better than this due to the higher redshifts of its source population.]

The LSST and SNAP approaches are thus quite complementary: LSST offers more raw statistical power, and SNAP will offer greater systematic-error control for shear and photo-z data. Large-scale WL power spectra will only be possible from the ground until such time as full-sky optical surveys are practical from space, perhaps sometime late in the next decade. The utility of LSST relative to SNAP will depend critically upon the limits

of our ability to control systematic errors from the ground. The precise division of responsibility between ground and space observatories is therefore still open for further study.

The current Pan-STARRS design, with four 1.8m telescopes, can undertake a shallower wide area weak lens survey, if programmatic considerations allow it to do so. Such a survey is likely to be state of the art in terms of areal coverage by 2010. Crudely speaking we might expect the ability of Pan-STARRS to constrain cosmological parameters to scale roughly as its throughput relative to the LSST design, assuming similar angular resolution for the two facilities, and further assuming an equal number of nights devoted to WL-compatible observations.

7.8. The Work Ahead

We've seen that weak lensing science puts tight constraints on the image quality of the LSST. Given a specific design for the telescope optics and an error budget for alignment errors, wind shake, surface tolerance, and so on, detailed models are needed to assess the effect on weak lensing statistics. For example, for the 8.4m LSST design, Monte-Carlo analysis of the effects of perturbing all degrees of freedom (rigid body position of the three mirrors plus their surface bending modes) should be carried out, including effects of AO system servo noise. Once these simulations are done, a realistic end-to-end Monte-Carlo simulation of the atmosphere and telescope optics needs to be carried out, to examine the variations in delivered PSF and their effect on weak lens shear measurements of typical source galaxies.

Further theoretical work is needed on how to optimally extract cosmological information from the weak lensing observations. This is a field of very active research, with qualitatively new ideas appearing on astro-ph as we write this section. Moreover, the current theoretical estimates incorporate noise and systematics in an ad-hoc way. While theoretical innovations for utilizing weak lens shear data have not been exhausted, it is time to undertake full Monte-Carlo simulations of the data pipeline including all known systematics and noise sources. Combining the instrument simulator, the observation simulator (including weather model), and the pipeline data simulator is a goal.

Bernstein, G. & Jain, B. 2003, astro-ph/0309332
Bernstein, G., & Jarvis, M. 2002, AJ, 123, 583
Browne, I. W. A. et al. 2003, MNRAS, 341, 13
Colley, W. N., Tyson, J. A. & Turner, E. L. 1996, ApJL, 461, L83
Firmani, C., D'Onghia, E., Avila-Reese, V., Chincarini, G., & Hernandez,X. 2000, MNRAS, 315, L29
Fukugita, M. & Turner, E. L. 1991, MNRAS, 253, 99
Fukushige, T., Kavai, A., and Makino, J. 2003, astro-ph/0306203
Hennawi, J. et al. 2004, in preparation
Hu, W. & Keeton, C. 2002 Phys. Rev. D 66, 063506
Hu, W., and Jain, B., 2003, astro-ph/0312395
Inada, N. et al. 2003, astro-ph/0312427
Jain, B., & Taylor, A. 2004, in preparation
Moore, B., Quinn, T., Governato, F., Stadel, J., & Lake, G. 1999, MNRAS,310, 1147
Oguri, M. et al. 2003, astro-ph/0312429
Refregier, A., & Bacon, D. 2003, MNRAS, 338, 48
Refregier, A. 2003, ARA&A, 41, 645
Riess, A.G. et al. 2003, ApJL, in press (astro-ph/0308185)
Sand, D.J., Treu, T., and Ellis, R.S. 2003, astro-ph/0310703
Seo, H.-J. & Eisenstein, D.J. 2003, ApJ, 598, 720

Song, Y.S., & Knox, L. 2003, astro-ph/0312175
Spergel, D. and Steinhardt, P. 2000, Phys. Rev. Lett. 84, 3760
Spergel, D.N. et al. 2003, ApJS, 148, 175
Takada, M. & Jain, B. 2003, MNRAS, in press (astro-ph/0310125)
Tyson, J.A., Kochanski, G.P., and Dell'Antonio, I.P. 1998, ApJ, 498,L107
Tyson, J.A. et al. 2003, Proc. Nuc. Phys. B 124, 21 astro-ph/0209632
Wittman, D. et al. 2001, ApJ, 557, L89
Wittman, D. et al. 2002, Proc. SPIE, 4836, 21
Wittman, D. et al. 2003, ApJ, 597, 218

8. Supernovae and the LSST

It is hard to overstate the impact of supernova research on cosmology in the past decade. Supernovae are the best standard candles at large distances (Gibson et al. 2000, Parodi et al. 2000). Supernovae provided the first of the triad of observational constraints on which the now-standard dark energy-dominated model of cosmology is based (Riess et al. 1998, Perlmutter et al. 1999). The challenge of the next decade of supernova research is to explore the physics both of supernovae themselves, and the nature of the redshift-distance relation. Massive surveys of supernovae at all redshifts with superb data are required for these goals.

The missions which will do this are the LSST and, presumably, the NASA/DoE Joint Dark Energy Mission (JDEM). Here we discuss two distinct supernova searches with LSST, one based on close to the standard cadence (Appendix A) to $z \approx 0.7$ (§ 8.1), and another, deeper survey using longer exposures probing the $z \approx 1$ regime (§ 8.3).

8.1. Moderate Redshift Type Ia Supernovae

8.1.1. Constraining the Dark Matter Equation of State

Using SN search simulation software written by J. Tonry, we have estimated the ability of the LSST to discover supernovae in its normal operating mode, which provides frequent all-sky coverage in multiple bandpasses. With a standard 250 $m^2 deg^2$ étendue, the LSST will discover and follow roughly 250,000 type Ia supernovae per year (down to a signal to noise of ~ 10) with a mean redshift of $z \sim 0.45$ and a maximum redshift of $z \sim 0.7$ (Fig. 14). The light curves from this standard observing mode will be marginally sampled both in time (roughly every five days) and in color (with only sparse coverage in three of the four bands). However, with recent and future advances in SN template fitting techniques, this sampling should provide adequate information for significant supernova studies. Assuming we can limit both the systematics and uncertainties due to lack of spectroscopic followup on all of the supernovae (an area of intensive current research by Barris et al., Prieto et al., both private communication, and others), the resulting redshift-distance relation data may be able to constrain w in the nearby universe to better than 1% when combined with priors from other experiments such as Planck. The degeneracy line of the weak lens technique in $w - \Omega_m$ space is nearly orthogonal to that of SN-based measurements of luminosity distance versus redshift, making these two measurements extremely complementary.

Supernova color statistics and good light curves, combined with a relatively small number of sample spectra, will allow us to test for any dependence of the supernova standard candle relation on parameters other than light-curve shape and extinction, shedding light on any systematic errors in the type Ia SN technique. One important component of the analysis, K-corrections, can be refined by "bootstrapping" using photometric redshifts of the host galaxies derived from the underlying LSST survey dataset.

The all-sky nature of the sample of 2.5 million type Ia supernovae that LSST will identify in its 10 years of operations will allow us to look for an angular dependence in the redshift-distance relation, thus determining whether the dark energy equation of state as characterized by w, and possibly even w', are directionally dependent. Any such signature would surely be an indication of fundamental new physics. Indeed, rough models show that $\sim 10,000$ supernovae are required to measure w to 1% with no priors (compare with the constraints that weak lensing will give in § 7), so over 10 years we should have ~ 250 independent samples of w to 1%. This of course assumes that there is no systematic floor in the supernova distance determination. However, with such a large sample, we could imagine deriving w independently for subsamples of supernovae with identical properties (e.g., light curve decay times, host galaxy types, etc.), to look for such systematics. Each subclass will provide an independent estimate of w, and consistency will indicate lack of serious systematic effects such as supernova

evolution.

With weak lensing (§ 7), constraining cosmological parameters requires a model for the growth of structure with epoch. By contrast, supernovae constrain cosmology by directly measuring the redshift-distance relation, and therefore the metric itself. If dark energy is a manifestation of something radically new in spacetime gravity, a comparison of the two approaches will reveal discrepancies which will give us clues about this new physics.

In addition, Hubble flows derived from SN Ia luminosity distances could be cross-correlated with CMB and large-scale structure data, yielding a separate constraint on dark energy with much higher S/N than those from studies of the integrated Sachs-Wolfe effect (Afshordi et al. 2004), complementing constraints from lensing.

8.1.2. Supernovae and Gravitational Lensing

The density of supernovae on the sky is such that we will be able to detect the magnification bias due to lensing from individual clusters. Given that we will have precise photometric (and in some cases spectroscopic) redshifts for these supernovae, and will accurately know the intrinsic (unmagnified) brightnesses of these objects, we will be able to put detailed constraints on models of the cluster mass distribution. Closer to the cluster core, supernovae will be gravitationally split; we estimate that one in a thousand high-redshift supernovae will be strongly lensed by a cluster or galaxy. Over the life of the survey, we will discover over 100 SN Ia's with multiple images and accurately measured time delays.

These lensing events have a big advantage over strongly lensed QSOs as the intrinsic brightness of the source is well known, thus allowing much stronger constraints on the lensing geometry. Well-modeled, strong lenses provide important constraints on the evolution of galaxy and cluster matter distributions, as well as H_0, itself an important ingredient in measuring the cosmic equation of state.

In addition, supernovae can be used as a generic weak lensing probe, complementing the techniques discussed in § 7. Indeed, weak lensing effects have been observed in supernovae (Williams & Song 2004); the LSST measurement of the galaxy distribution will allow such analyses to be done with great precision.

8.1.3. The Physics of Supernovae

Supernova light-curves will be followed for many years as the supernova ejecta expand to reveal the inner workings of the explosions, leading to new understanding of the SN Ia mechanism through, e.g., the interaction with the circumstellar environment. Very late-time light-curves can also provide additional constraints on host-galaxy reddening using the color constancy of their pure iron recombination spectra.

Core-collapse supernovae provide the instantaneous star formation rate, and LSST will accurately map this rate by finding many tens of thousands of these events in the redshift range where it is known that the cosmic star formation rate has changed most dramatically. Accurate star formation rates will be calibrated as a function of galaxy type and (given spectroscopic follow-up) emission line flux.

The LSST data will enable an accurate determination of the luminosity function of type Ia supernovae, which will help constrain progenitor models and explain the observed dependence of the luminosity distribution of these supernovae with galaxy morphology.

Core-collapse supernovae are now believed to be the source of gamma-ray bursts, but the mechanism producing the bursts remains a mystery. Some models predict a delay between the supernovae explosion and the formation of the gamma-ray burst, and afterglows provide a variety of information on the gamma-ray emission mechanism itself. Testing these ideas requires monitoring large areas of the sky to see if a supernova has exploded days, weeks,

or months earlier (§ 5). The LSST enables natural tests for such models, especially if the delay timescales are longer than a few days.

8.2. Moderate Redshift Type II Supernovae

LSST will also discover nearly as many Type II supernovae as Type Ia's, again with finely-sampled light-curves in many colors. Type II supernovae can be used as complementary distance indicators (the EPM method; Schmidt et al. 1994, Hamuy et al. 2001) to the type Ia's, although to smaller redshifts due their fainter intrinsic luminosities. Note also that the EPM method requires extensive synoptic spectroscopy in order to determine the photospheric expansion velocity as a function of time, thus such work will require extensive access to 8-10 meter class telescopes with spectrographs. Most type II supernovae can be distinguished from other types of SN by the duration and color evolution of their light-curves. The supernova rates themselves, together with the photometric redshifts which LSST will obtain of their host galaxies, will be a direct measure of the star formation history of the universe. Late-time light-curves will provide a direct measure of type II supernova ^{56}Ni (and hence iron) yields. The amount of iron which is released in the supernova explosion depends sensitively on the fraction of the total produced by explosive burning in the silicon shell falls back into the compact object at the center. The watershed mass coordinate dividing what falls back and what escapes (the so-called "mass cut") can be measured from the ^{56}Ni yield, and is crucial for our understanding of cosmic chemical evolution of iron-group elements, and the mass function of compact remnants.

8.3. High Redshift Supernovae

As we have seen, LSST will do a definitive job of finding and obtaining light-curves for supernovae with $z < 0.7$, where the dynamical effects of dark energy, especially evolving values of w, are maximized. The space-based JDEM project (SNAP or DESTINY, for example) will focus on higher redshifts over much smaller areas on the sky. Such a satellite is expected to measure Ω_m to 0.03, w to 0.07, and w' to 0.3 employing no additional information (Kim et al. 2004).

The LSST also has the potential to probe these higher redshifts. At $z \approx 1$, supernovae are red, and fainter than can be followed in single exposures reaching $r \sim 24$. Thus this second, deeper LSST supernova sample will come from a "staring" mode search of a more limited area of sky, with longer exposure times and frequent repeat exposures on the chosen fields. Because LSST surveys such a large volume per field, and supernovae are relatively frequent events, a relatively few fields spaced about the sky would suffice to measure cosmological parameters as a function of redshift and direction. Observations carried out over a single night to $V = 26$, $z = 25$ (10–20 minutes per band) are required every five days in each of four bands (not necessarily on the same day). Individual fields would be followed for four or five months to allow for complete coverage of detected supernovae light curves (the length that a single light curve could be followed is roughly three months at $z = 1$, but with a three-month campaign, all supernovae that started after the first few weeks of the period would have only partial coverage). If we assume no evolution in the supernova rate, we estimate that this survey will yield more than 1,400 supernovae in a 3° field in a dedicated year-long campain. These light curves will have unprecedented detail, with 60-100 photometric points per light curve total in five bands. The sample will have a mean redshift of 0.75 and extend (with more limited color coverage) beyond $z \sim 1.4$. Such detailed color information will allow fitting for photometric redshifts from the supernovae themselves. Realistic simulations (see Fig. 15) show that fitting multi-color light curves of the quality we hope to get for host-galaxy reddening, redshift, and distance modulus simultaneously yields redshift errors of $0.7\%(1+z)$ and extinctions accurate to 0.03 in A_V.

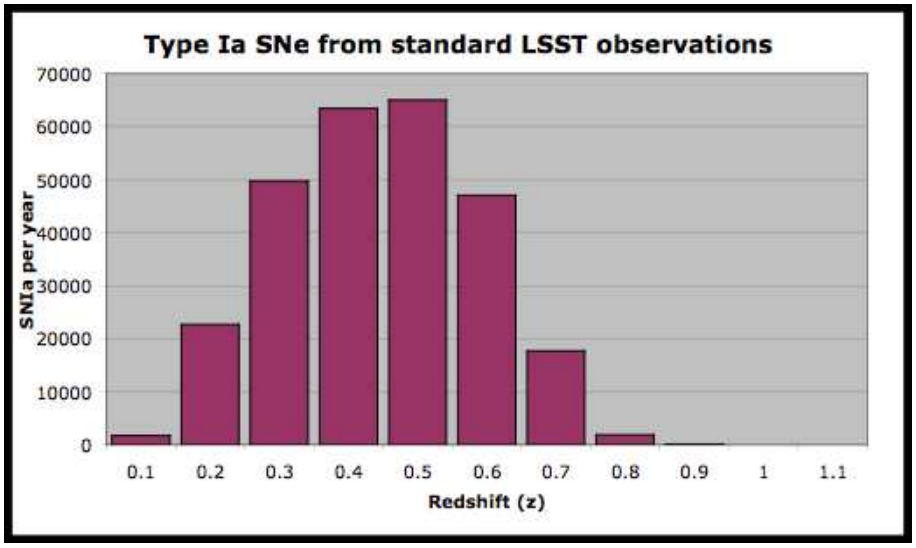

Fig. 14.— Distribution with redshift of the supernova discovery rate per year from a search carried out as part of the normal operating cadence.

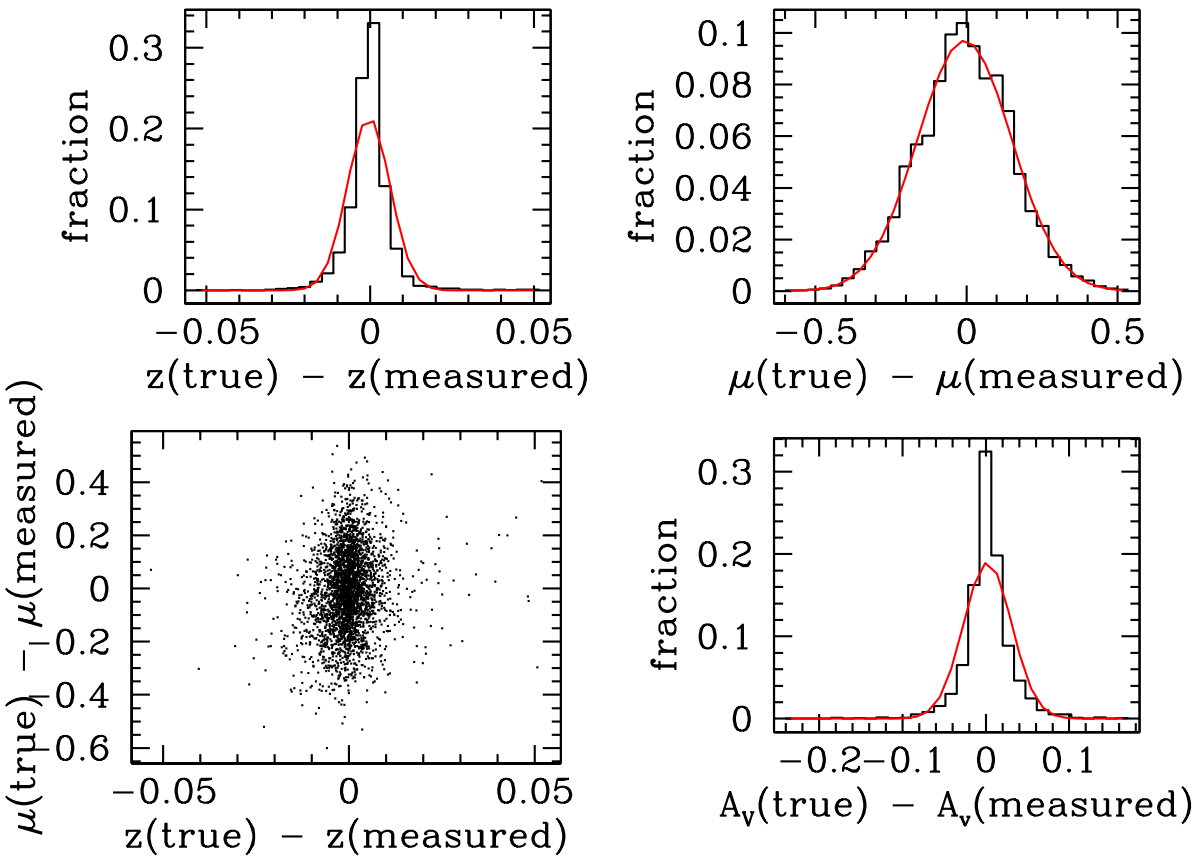

Fig. 15.— Errors in recovering $z < 1$ SN properties from light curves obtained in the "staring mode" described in § 8.3. (clockwise from top left): a) redshift errors; mean = −0.0004, σ_z=0.007. b) distance modulus μ error; mean = −0.006, σ_μ=0.16. c) reddening error; mean = 0.002, σ_{A_V} = 0.03. (d) correlation between redshift and distance modulus errors. The covariance is weak.

The photometric redshifts of the supernovae can be compared to those of their host galaxies, and since the area on the sky is limited, it is reasonable to assume that spectroscopic followup can be obtained for many objects. Figure 16 shows a typical light-curve for a supernova at $z = 0.8$ from a complete simulation of the high-redshift supernova sample (Pinto 2004, private communication) incorporating a realistic observing cadence, telescope and detector throughput, and the uncertainties inherent in the light curve fitting procedure. Using a randomly selected sample of 2000 supernovae from one realization, Figures 17, 18, and 19 show the resulting constraints on cosmological parameters possible in the first year.

Note that this program may naturally mesh with the deep Kuiper Belt program described in § 4, although note that the KBO program requires returning to a given field with a much slower cadence than the needed to follow supernova light curves.

8.4. Requirements

To employ LSST as a true supernova factory we need to give careful thought to design requirements for the LSST survey(s). A survey with five filters, returning to the same field every three or four nights as specified in the "universal cadence" (Appendix A), would cover one-quarter of the sky per cycle. As the year progresses, new fields in the East would replace older fields in the West, eventually covering most of the visible sky.

A deeper survey reaching $V = 26$ and $z = 25$, however, requires 10-20 minutes of accumulated exposure in each filter, with revisits of five to seven days; visits in the different filters need not be done on the same day. This change in the exposure time from the standard cadence results in a reduction of N_{FOV} by factors of 30-50. It is therefore necessary to define some fraction of every night which is devoted to this long exposure cadence. If that fraction were 10%, such a program would discover and study 30000 type Ia supernovae per year out to a redshift $z = 1.2$ and a comparable number of core-collapse supernovae, typically to $z = 0.6$ but with significant numbers to $z > 1$ (for example, some progenitors of gamma-ray bursts such as SN 1998bw). BVRIz light-curves (the Y filter would be even better with a very red-sensitive detector) would be obtained in parallel with the search and provide a means of discriminating between supernova types (with spectroscopic followup of some small fraction of events to determine reliability). The combined data sets will provide photometric redshifts to host galaxies. Work remains to be done to merge this program with the suggested deep KBO program described in § 4; as pointed out before, the two need somewhat different filters (the KBO science really only needs observations in two filters), and different cadences.

We touch briefly on other requirements.

Photometric Calibration: Absolute photometric calibration to 2% is adequate for supernova science. The calibration of the zero points between filters (the absolute spectrum of Vega is not known to 1%) is a serious issue bearing on the subject of K corrections. The stability of the filter response functions over time (or the frequency of the measurement of the throughput of the system as a function of wavelength) is also an issue.

Sky Coverage: To reach the goal of well-sampled light curves for hundreds of thousands of supernovae requires observations of the high-latitude sky several times per month to $R \sim 24$, similar to what we've seen in other chapters.

Seeing: Good seeing is important for supernova work, to go faint, but more importantly, to separate supernovae photometrically from their host galaxies. This is especially important at high redshift, where the experience with HST is that the photometry is *far* superior to what is obtainable from the ground. Of course, properly sampled images are important as well. Accurate PSF photometry on a complicated background requires an accurate model for the PSF, which requires that the PSF be relatively stable as a function of position on the focal plane. There are no strong requirements on astrometric accuracy. Like the variability searches, rapid reduction of the data is

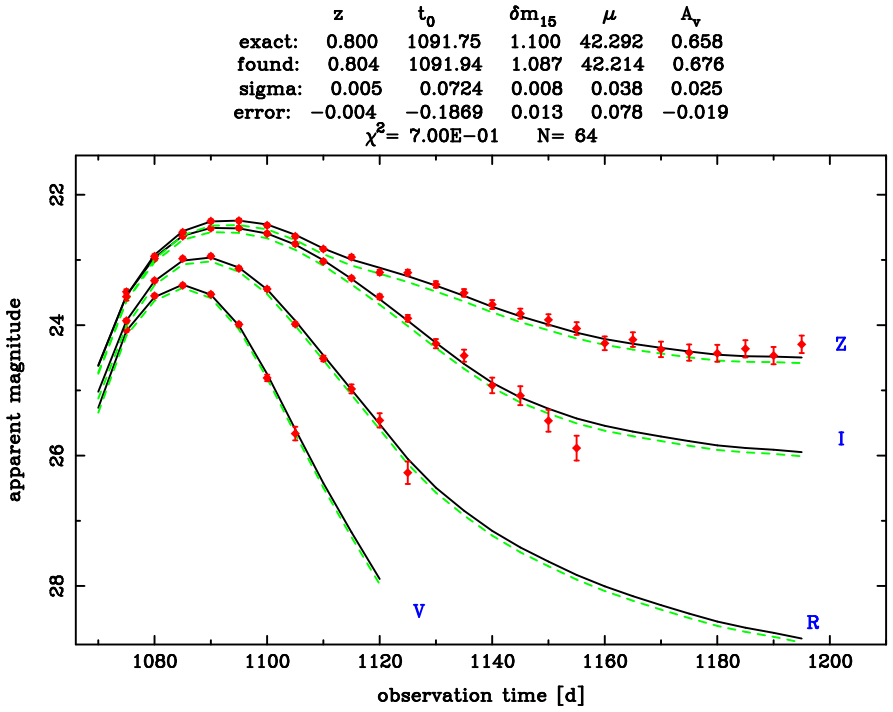

Fig. 16.— Typical light-curve and fit resulting from observation of a $z = 0.8$ supernova observed with 10 minutes of exposure time in VRI and 20 minutes in z every five days.

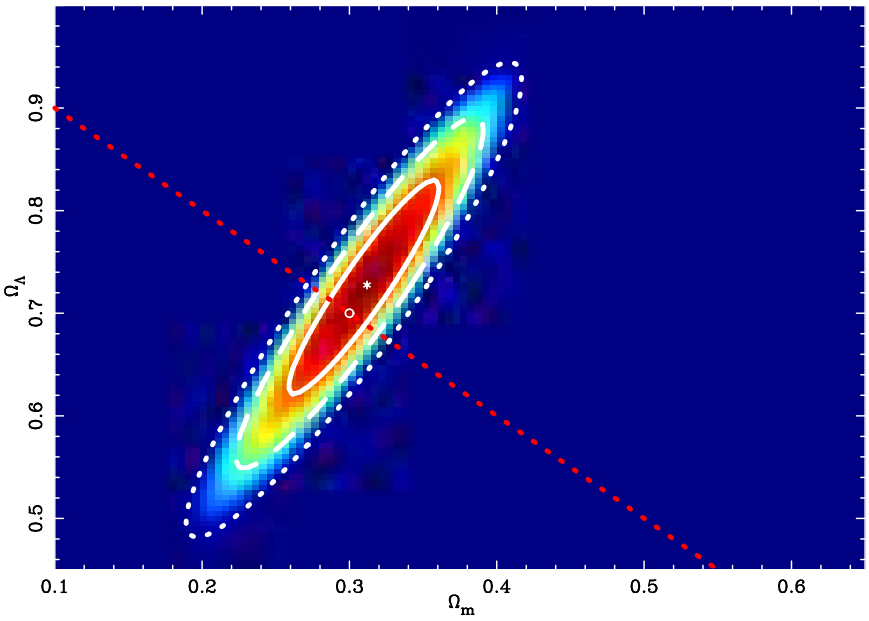

Fig. 17.— 1, 2, and 3 σ confidence contours for $(\Omega_M, \Omega_\Lambda)$ for 2000 supernovae to a redshift of $z = 1$ from one realization of the simulation described in the text. No prior knowledge of H_0 is assumed, and w was assumed to evolve with redshift as $w = w_0 + w_1(1 + z)$.

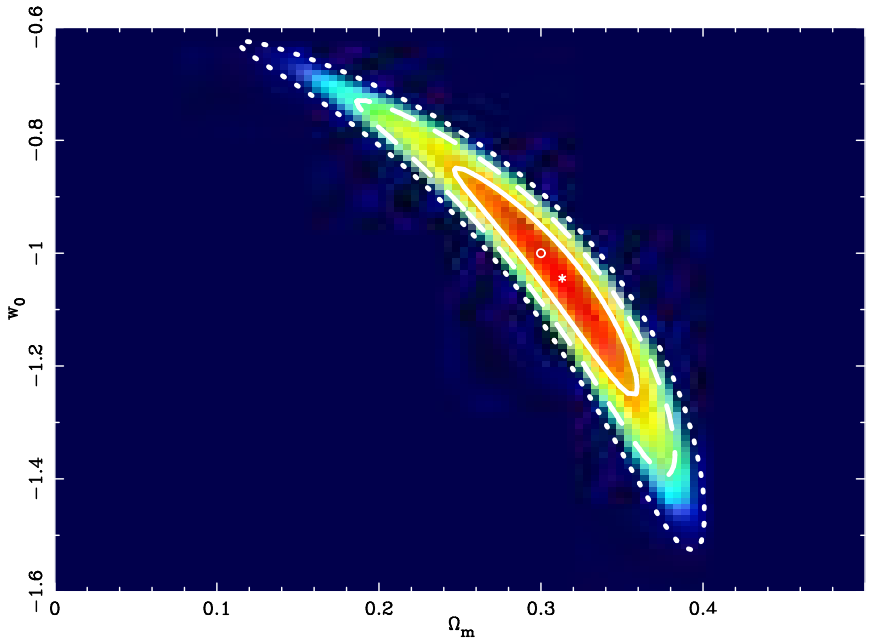

Fig. 18.— 1, 2, and 3 σ confidence contours for Ω_M and w (here assumed constant) for the supernova sample used in Figure 17. No prior knowledge of H_0 was assumed.

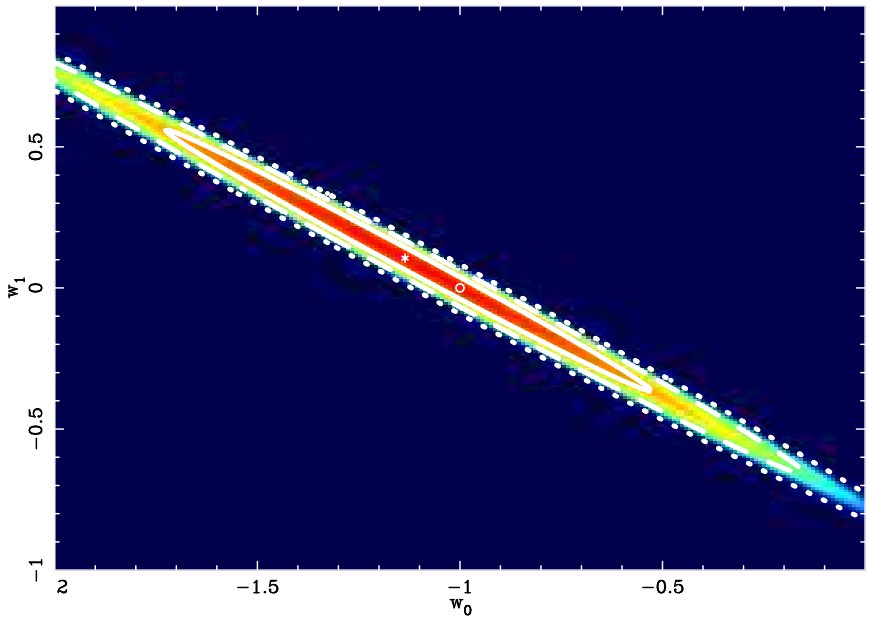

Fig. 19.— 1, 2, and 3 σ confidence contours for (w_0, w_1) $(w = w_0 + w_1(1+z))$ for the supernova sample used in Figure 17. No prior knowledge of H_0; $\Omega_{tot} = 1$ was assumed.

important to allow any follow-up work to be done.

Auxiliary data: We have made reference here to spectroscopic observations of supernovae obtained on other facilities. The full scope of what is needed here remains to be quantified; however, many of the core science goals described in this chapter can be achieved with just enough spectroscopic data to properly calibrate and test supernova photometric redshift techniques. Indeed, the simulations described above assumed that redshifts were determined from supernova light curves, not from spectroscopy.

How much of supernova science is likely to be done by the time LSST comes on line? The largest current high-redshift supernova surveys are the ESSENCE and CFHT Legacy surveys. ESSENCE is designed to obtain 230 supernovae with $0.2 < z < 0.75$ to measure w to 10%. Optical observations are taken with the CTIO 4m and zJ photometry and spectroscopy are obtained using various 8m class telescopes. The CFHT Legacy Survey is expected to detect about 700 type Ia supernovae over a five-year period with a similar program, and also plans to be able to determine w to 10%. Pan-STARRS will discover of order 10^4 supernovae per year (Tonry et al. 2003). As we have seen, LSST will obtain orders of magnitude more supernovae, and determine w to much higher precision. It is already known, from WMAP and other analyses, that w is within 20% of -1 (the Cosmological Constant value). Understanding the physical nature of dark energy requires measuring any deviation of w from -1, and its time and directional dependence, to superb precision.

Afshordi, N., Loh, Y.-S., & Strauss, M.A. 2004, PRD, 69, 083524
Gibson, B.K. et al. 2000, ApJ, 529, 723
Hamuy, M. et al. 2001, ApJ 558, 615
Kim, A.G., Linder, E.V., Miquel, R., & Mostek, N. 2004, MNRAS 347, 909
Parodi, B.R. et al. ApJ, 540, 634
Perlmutter,S. et al. 1999, ApJ, 517, 565
Riess, A.G. et al. 1998, AJ, 116, 1009
Schmidt, B.P. et al. 1994, ApJ 432, 42
Tonry, J. et al. 2003, ApJ, 594, 1
Williams, L.L.R., & Song, J. 2004, MNRAS, in press (astro-ph/0403680)

9. Other Science Topics

The previous sections describe the principal scientific drivers for the LSST. However, the LSST data set will allow a much broader range of science to be carried out. Here we briefly touch upon a few scientific topics, although as emphasized in the Introduction, perhaps the most exciting science that LSST will accomplish are things that we have not even anticipated at this point.

9.1. Main-belt Asteroids

Current estimates of the number-magnitude relation for main belt asteroids are summarized in Figure 20. The LSST will determine their orbital parameters, determine accurate photometry in a variety of bands, and give detailed information on their variability. Their orbital distribution and surface properties contain a detailed record of the dynamical and chemical history of the solar system.

Accurate colors of asteroids are particularly interesting. The SDSS found a strong segregation of asteroid dynamical families in the space spanned by SDSS colors (Ivezić et al. 2002). Asteroid dynamical families are clusters of asteroids in orbital element space (Hirayama 1918; for a review see Binzel 1993). Hirayama proposed that the families may be the remnants of parent bodies that broke into fragments. Figure 21 shows the main-belt asteroid distribution in the space spanned by proper semi-major axis and the sine of the orbital inclination angle, with the points color-coded according to their colors measured by SDSS. The two main asteroid taxonomic classes, S (silicate) and C (carbonaceous), correspond to red and blue shades, respectively. The green shade is dominated by the Vesta family. The families are apparent as clumps in this figure. It is striking how homogeneous and distinct the colors are within each family.

Szabo et al. (2003) suggest that this palette of colors is due to space weathering – the bombardment of asteroid surfaces by micrometeorites, cosmic rays, solar wind and UV radiation that alters the chemistry of the surface material (Zeller & Rouca 1967).

The LSST would improve the currently available data in several ways:

1. A factor of ten increase in the number of detected objects with accurate colors;
2. A Smaller size limit by about a factor of three.
3. Orbital parameters will be determined for practically all objects, thereby increasing the sample of objects with both orbital parameters and colors by over a factor of twenty.
4. The variability information, which carries important information about the physical state of an asteroid (e.g., solid body vs. a rubble pile) will be available for practically all objects, thereby increasing the sample size by over a factor of hundred. These new data will constrain the size-strength relationship, which is a fundamental quantity that drives the collisional evolution of the asteroid belt.

9.2. Quasars and Active Galactic Nuclei

Quasars are the manifestation of material accreting into supermassive black holes (SMBH) at the cores of galaxies. The ubiquitousness of SMBH in galaxies, and their tight correlation with properties of their host galaxies, strongly suggest that quasar activity and/or black hole growth are an important phase in the lives of essentially all massive galaxies. Understanding this connection requires large and well-calibrated samples of galaxies and quasars to high redshift. The LSST will generate a dataset to do exactly this.

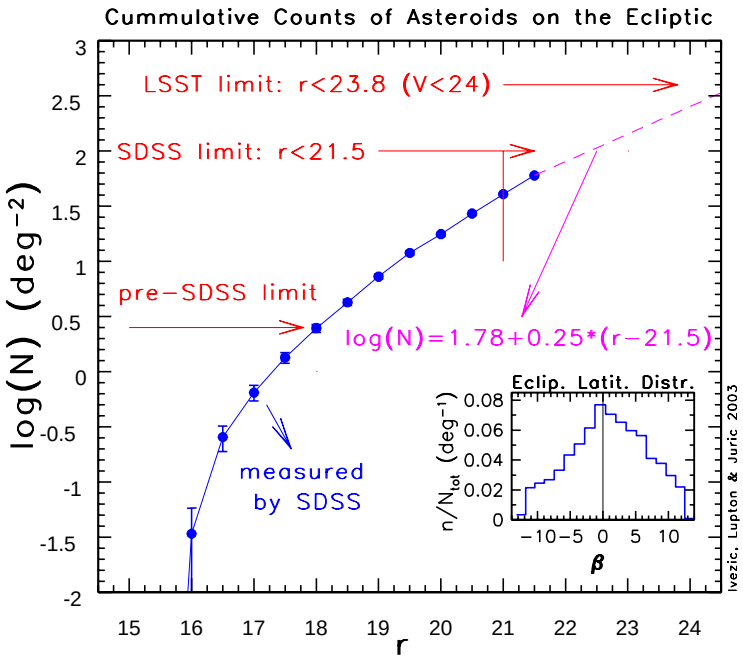

Fig. 20.— The cumulative counts of main belt asteroids on the Ecliptic, as measured by the SDSS for $r < 21.5$ (Ivezić et al. 2001), and extrapolated to $r = 24$ (this extrapolation is supported by recent Subaru survey of 3 deg^2 which finds a density of 290 deg^{-2} for $R < 24.4$, Yoshida et al. 2003). Taking into account the ecliptic latitude distribution, shown in the insert, the effective sky area is ∼6,000 deg^2, and implies that LSST will discover, and determine orbital parameters for over a million main-belt asteroids.

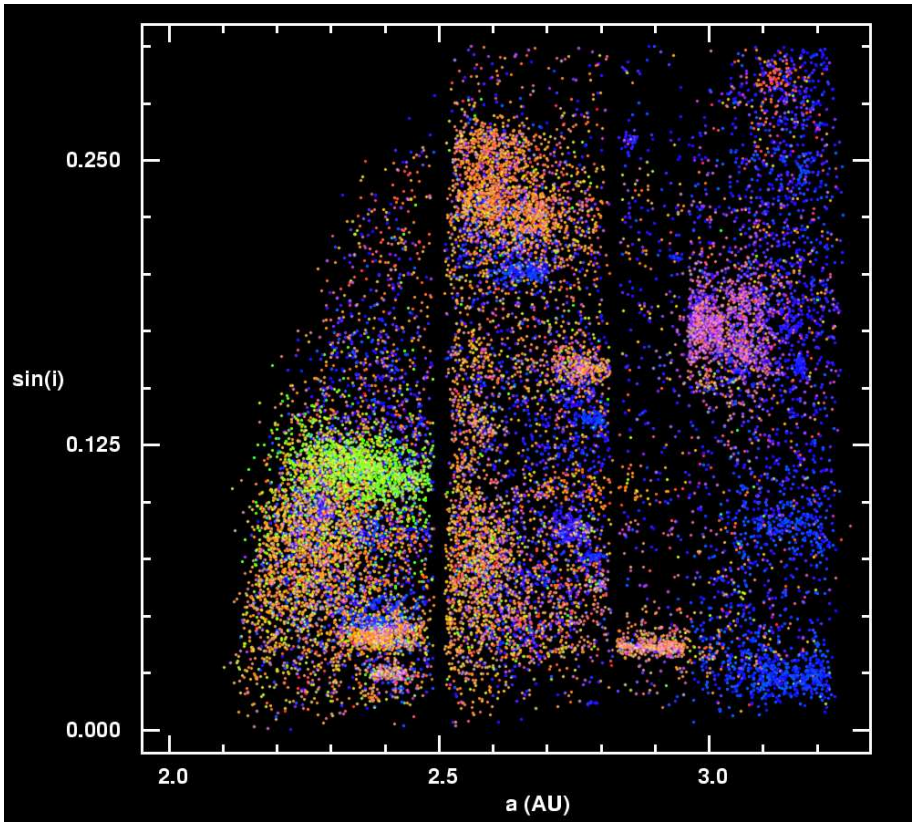

Fig. 21.— The distribution of 27,000 asteroids with available proper orbital elements and SDSS colors in the space spanned by the proper inclination and semi-major axis (approximately, the x axis is proportional to the distance from the Sun, and y axis is proportional to the distance from the orbital plane). The dots are colored according to their measured SDSS colors (Ivezić et al. 2002). The clusters of points are asteroid dynamical families. Note their strong color segregation, suggesting a common origin.

The largest quasar samples to date, those of the SDSS and 2dF redshift surveys, were selected from their distinctive colors relative to those of stars (see Richards et al. 2002). In particular, $z < 2.5$ quasars which dominate the number counts, have bluer $u - g$ (or $U - B$) colors than do ordinary stars. It remains unclear whether the LSST will include a U filter. However, it will gain exquisite information on the variability of every object. Quasars are known to vary in the optical on timescales from weeks (shorter for blazars) to a year, depending on luminosity; information on the timescale and amplitude of variability, together with colors, will make selecting quasars quite straightforward. If the colors are measured over timescales less than the variability, these colors can be used to derive photometric redshifts for quasars, which will be adequate for determining trends of their properties with cosmic epoch. Moreover, the variability studies will allow the isolation of quasars without "typical" quasar colors, including populations of quasars with luminosities comparable to that of their host galaxies.

Quasars grow both by steady accretion of material through a disk, and presumably through the sudden disruption of stars that are tidally shredded by the black hole. This type of event is very rare, but has the potential to teach us a great deal about the feeding of black holes. The LSST quasar sample and associated variability will be large enough to look for, and trigger on, sudden flares, which could then be followed up with other telescopes.

The LSST should also be an effective discoverer of the most distant quasars. All known quasars with redshift $z > 5.5$ have been discovered as i-band dropouts in the SDSS; the presence of a Gunn-Peterson absorption trough in those quasars at $z > 6.2$ indicates that we are probing the end of the epoch of reionization. Our understanding of this important stage in the thermal history of the universe is very much limited by the small number of objects known, and therefore the small number of sightlines probed through the intergalactic medium at these very high redshifts. Such quasars are very rare (the SDSS has discovered only 3 at $z > 6.2$ over roughly 5000 deg^2), thus a very large angle survey is needed with accurate deep photometry in red bands. The LSST will go four or five magnitudes deeper in its stacked data than the SDSS, and if it includes a Y filter, will probe to $z \approx 7.6$. The quasar luminosity function at high redshift is likely very steep at the bright end, thus the LSST has the potential to discover hundreds or even thousands of quasars with $z > 6$, opening avenues for statistical analyses. The mere existence of quasars, and the $10^9\,M_\odot$ black holes which power them, is difficult to understand so soon after the near-perfect homogeneity which held at the epoch of last scattering; having a large statistical sample will sharpen this question. And studying the evolution and spatial structure of the ionization state of the IGM at the end of reionization will constrain the energy output from the first light sources.

9.3. LSST and Large-Scale Structure

With deep, accurate multi-color data, the LSST will allow accurate photometric redshifts to be determined for billions of galaxies to redshifts of 3 and perhaps higher. Galaxy clustering can be measured as a function of galaxy luminosity and type from the present to $z = 1.5$ and perhaps beyond. Existing redshift surveys using ten-meter-class telescopes have already shown evidence for strong, biased galaxy clustering at $z \sim 3$ on relatively small scales, but measuring the galaxy power spectrum on large scales at high redshift is something that the LSST can do well. Acoustic oscillations in the photon-baryon fluid after matter-radiation equality imprint a series of "baryon wiggles" in the power spectrum in large scales; these also cause the dramatic peaks seen in the power spectrum of CMB fluctuations. These are difficult to measure at low redshifts, as non-linear growth of the density fluctuations tends to smear them out, but at high redshift, non-linear effects should be much lessened. The amplitude of these wiggles depends on the ratio of Ω_b to the total mass density, while the wavelength scale on which these wiggles are seen is dependent on the geometry of the universe and the Hubble Constant. Thus measurement of these predicted wiggles allows a crucial consistency check with the results of WMAP (and even more basically, as they

are required by our cosmological models, their presence would be a reassurance that these models are correct), and would allow an independent determination of the Hubble Constant.

Binzel, R.P. 1993, ASP Conference Series, Vol. 63, 251

Hirayama, K. 1918, AJ, 31, 185

Ivezić, Ž. et al. 2001, AJ, 122, 2749

Ivezić, Ž. et al. 2002, AJ, 124, 2943

Richards, G.T. et al. 2002, AJ, 123, 2945

Szabo, G. et al. 2003, astro-ph/0309112

Yoshida, F. et al. 2003, PASJ, 55, 701

Zeller, E.J., Rouce, L.B., 1967, Icarus, 7, 372

10. Realizing the LSST Multiplex Advantage

The preceding sections set forth a strong scientific justification for the LSST system. As described in § 3, 10-second exposures that reach to 24th magnitude are desirable for a complete sample of faint potentially hazardous objects (PHAs), mandating an effective collecting area equivalent to an aperture of greater than 6 meters. The sky coverage rate required in order to detect 80-90% of the PHAs of greater than a few hundred meters in diameter requires an étendue of 250 $m^2\,deg^2$. This survey could be carried out in a single broad optical passband.

Achieving the requisite surface density of background galaxies for weak lensing studies of the nature of dark energy (§ 7) requires the ability to detect galaxies down to $r_{AB} = 26.5$ at 10σ, with a high degree of control of systematics in the PSF shape. Although the shape measurements can be drawn from a single optical passband, disentangling the source and lensing galaxy populations, and studying the shear as a function of look-back time, require deep images in at least five bands in order to obtain robust photometric redshifts. Achieving the sky coverage, the depth, and the wavelength coverage needed for weak lensing thus similarly demands an étendue in excess of 270 $m^2\,deg^2$.

A deep survey for Kuiper Belt Objects over an appreciable fraction of the ecliptic plane requires deep ($r \approx 26$) large area images with several repeat exposures; the science goals described in § 4 again suggest an étendue of order 250 $m^2\,deg^2$.

A well-planned observing cadence (Appendix A) will provide good light curve coverage over the sky, allowing studies of astrophysical variability (§ 5) on timescales from tens of seconds to years. Going to faint magnitudes in multiple passbands over large areas of sky will open as-yet unexplored regions of parameter space. For a fixed amount of telescope time there are tradeoffs between sky coverage, temporal sampling density, and wavelength diversity. Certain science goals, such as a multiband deep survey for supernovae (§ 8) have very specific requirements on these tradeoffs, given that LSST will not just be discovering supernovae, but generating their lightcurves as well.

For each of the science goals outlined in this document, we have given a sense of the optimal observing strategy, cadence, and filter choice. While the choices of each goal are similar, they are not identical. Thus the question arises, to what extent can the LSST merge these different observing programs into a single coherent strategy such that the images obtained can be used for multiple programs? Such a synergy potentially provides tremendous leverage over the traditional approach to astronomical observations, whereby large aperture telescope time and the resultant data are awarded on a project-by-project basis, for each group to work on independently. The LSST is a survey project, like, e.g., SDSS and 2MASS, which use dedicated telescopes to produce a set of images and object catalogs to be used for a broad range of scientific projects. On the other hand, time domain surveys such as MACHO, OGLE, ROTSE and their kin have been carried out on 1m class telescopes, often optimized for a particular science objective. The LSST represents the first opportunity to carry out a multipurpose survey over the entire accessible sky, to faint flux limits, with rapid revisit rates.

Any merged observing program involves tradeoffs unless the optimal observing program is identical for all science goals. We assert that the LSST observing program *should* take a diverse set of goals into account, with appropriate weighting, and the observing program should be crafted to enable the maximum amount of cutting edge science. A clear example arises in considering passbands for NEO detection. While the solar system observations would benefit from using a single broad optical passband, if the rapid-revisit images were obtained in two different passbands the marginal gain for variability science is huge. Gamma ray bursts, supernovae and classification of generic variability all derive tremendous gains from the added information on the spectral energy distribution of the variable sources. The slight degradation in detection limits for NEOs is more than compensated by the scientific gains obtained from multiband images.

Studies of the outer regions of the solar system benefit from deeper images over a narrower field in two or more passbands, as do supernova observations. Obtaining dense temporal sampling of supernova light curves also requires more frequent revisits than are needed for PHA orbit determinations. This argues for allocating a segment of the LSST observing to sequences of multiband images, in regions of the sky that would benefit both the KBO and supernova studies. (It is worth emphasizing that neither KBO nor supernova science, unlike, e.g., weak lensing and NEO studies, require sky coverage of tens of thousands of square degrees.) This is an important ingredient in searching for serendipitous new sources of variability as well.

It is difficult to arrive at a closed-form analytic optimization of LSST observations that will meet a diverse set of science goals. Computer simulations are a very powerful tool for comparing different scenarios, and number of these are under development in the community. These will prove to be important in appraising different proposed observing schemes. In the absence of a detailed comparison, and at the risk of misjudging scientific priorities once the LSST system in on the sky, we suspect that the optimal LSST observing program will likely include two or three distinct components. One example implementation might include:

- An initial four- or five-band survey of the entire accessible sky, at the depth achieved from a stack of perhaps 3-10 images, which could serve as an early data deliverable for the LSST system. The filters would be broad-band, and would cover the optical window from u or g to z or Y. These data would be absolutely crucial for setting the initial photometric and astrometric zeropoints for the LSST. This is also a sensible commissioning project for the LSST.

- Deep multiband imaging over a few hundred square degrees along the Ecliptic plane, to detect and monitor supernovae, KBOs, and other variables, using the cadence described in § 4. Note that the KBO program and the supernova program have somewhat different choices of filters and ideal cadences; this is definitely an area for further work. Of order 10% of total LSST time could be allocated to this program.

- An ongoing implementation of the Universal Cadence described in Appendix A. Cycling through filter pairs, with images taken at 15 minute separations, would allow the images to be stacked for photometric redshift, weak lensing, and astrometric studies. It is clear, however, that this requires further thought. As discussed in § 8.1, the ideal cadence for supernova work requires somewhat finer time resolution, and broader filter coverage, than the Universal Cadence naturally gives.

With this combined approach, the LSST system can realize the goal of feeding multiple and diverse science objectives from a common data stream. Adjusting the balance between and among observing programs will fall to the LSST scientific management structure, once the system is in operation. This group would presumably act to achieve the sort of balance that we currently look to Time Allocation Committees to provide for present era telescopes.

The scientific *value* of the LSST will be greatly increased with the adoption of multipurpose observing strategies of the sort described here. The larger the ètendue, the more natural it is for a single observing strategy to address multiple science goals. We encourage the cognizant scientific communities to remain flexible, and to consider the *overall* scientific merits of different proposed strategies in the context of their effectiveness for any specific scientific goal.

11. LSST Data Access: A recommendation

The LSST system will provide an unprecedented flow of astronomical information, with a premium on real-time exploitation of detected variability. Some of the LSST science can be achieved without additional "followup" observations, while other programs will greatly benefit from coordinated spectroscopy, for example. This will require rapid analysis of the LSST image stream, and prompt delivery of "alerts" to the community.

In addition to the real-time aspect of the LSST data set, the object catalogs that will be produced by the system are also of tremendous value to the community. This survey will go many magnitudes deeper than any previous survey with similar sky coverage. The multiplicative leverage of surveys in astronomy has a long history, from the original surveys using photographic plates, to the digital versions of these images, and more recently multicolor surveys such as SDSS and 2-Mass. We expect that the LSST object catalogs will provide a qualitative leap in scientific reach, for programs ranging in scale from the solar system to the cosmological.

In order to achieve the maximum scientific and educational productivity, the LSST data should be made accessible as promptly as possible. Our committee advocates community access to the LSST images and data products with no proprietary data period. Our enthusiasm for this principle is tempered by the recognition that there are three non-trivial issues that must be addressed in the data distribution policy:

1. For the data to be of genuine value to the community, it must be trustworthy. This requires a credible and timely Quality Assessment (QA) program, and

2. For a data set of this size, resource limitations will likely constrain the nature and scope of queries that can be supported, and finally

3. The LSST enterprise must decide what level of user support is going to be provided.

The question of QA is a perennial problem for astrophysical data sets. On the one hand, the community is frequently clamoring for prompt access to new data sets. Delays in data release are treated as self-serving attempts to skim off the cream science. On the other hand, projects that provide catalog or image data that retain instrumental artifacts, calibration non-uniformities, or other pathologies are derided as not having done a sufficiently thorough job, and the project's credibility is compromised. A further consideration is that pathologies in the data set only come to light once a scientific project is close to completion, often requiring adjustments to the software and a re-reduction of the pertinent images.

Determining the right balance between maturity of the data products and prompt access is beyond the scope of this report. We encourage the LSST proponents to engage the astronomical community in a discussion of how this might be addressed. We suspect that a mixed approach will be attractive, in which preliminary classifications of detected variability are released immediately, and "catalog" level data products are provided to the community as certain verification criteria are attained. In all cases, clearly defined criteria must be set forth and adhered to. We recognize that having a few key science projects as an integral part of the LSST system is a good way to achieve QA on the data set, and we encourage this approach.

Storage, bandwidth and computational resources will all limit access to the LSST data set. The astronomical community is accustomed to dealing with resource limitations, in particular telescope time. The approach adopted there is to award access on the basis of scientific merit, through a peer review process. We envision a similar approach for resource-intensive LSST data requests. There is likely to be a logarithmic distribution of demands on the LSST data archive, with many small requests ("Give me the light curve for this object") and a few very demanding ones ("Find all instances on the sky where detected proper motions are likely to produce a microlensing

event"). While the LSST system needs to accommodate both instances, we suspect the large-scale queries will prove the most stressful to the system. This is easily dealt with by instituting an analog of traditional Time Allocation Committees, although in this case it will be an LSST Transaction Allocation Committee.

The issue of user support has clear cost implications. Apart from the obvious desirability of having clear documentation and straightforward user interfaces, will the LSST system provide the equivalent of a help desk? There are two inter-related issues here. One challenge the LSST faces is to educate its user base in the exploitation of the data set. This can be achieved in part by holding workshops in conjunction with scientific and educational meetings, and by providing clear examples and FAQ documentation. Even with a vigorous program of user education, however, there will be a group of users who will request help in achieving ambitious goals. For the LSST project to achieve its full potential, we consider the staffing of a "help desk" as a very productive investment.

Our working group strongly endorses the idea of having no proprietary period for privileged data access to the LSST data. We recognize that giving data away is not a simple task, and we do think there are tractable approaches to the issues discussed above. We urge the NSF to encourage the LSST endeavor to adopt the philosophy of maximizing the scientific and educational benefits of the project by striving to provide maximum availability of the LSST data, working in close conjunction with the research and educational communities.

12. What lies ahead

As part of the deliberations of the LSST Science Working Group, we found ourselves reiterating some of the major hurdles, concerns, and areas for further work that any implementation of the LSST will have to address. The most important of these we list last, namely understanding how the scientific drivers for LSST, and the compelling nature of the project, will change as new surveys come on line. We do not offer solutions to these problems here, but simply list them and give a few comments for each.

Photometric calibration:

The science goals of the LSST push us to photometry accurate to 1-2% in multiple bands, free of systematics as a function of magnitude or position on the sky. The SDSS is our closest experience to this, and is approaching the 2% goal, but it has been quite a bit more difficult than imagined, due to a myriad of issues: flat-fielding, accurate determination of the PSF, developing an accurate grid of standard stars, uncertainties in filter response, etc., etc. The LSST problem will be made substantially easier because each area of sky will be imaged multiple times: this will allow many of the systematic effects to be averaged out. But this needs to be thought through very carefully in both the telescope design and cadence strategy. Similarly, the problem of calibrating data taken under non-photometric conditions is one that hasn't been tackled yet.

Indeed, we need a way to flag each data frame taken with the LSST with some measure of the photometricity of the given night. One way to do this is to use auxiliary instrumentation to determine the transparency of the atmosphere, such as a Differential Image Motion Monitor (DIMM), an all-sky camera sensitive to clouds (in the optical, or preferably at 10 microns), and/or a small optical telescope which measures standard stars all night, thereby determining extinction coefficients.

Note that the photometric calibration problem naturally breaks into two parts: relative and absolute. Relative calibration requires that the photometry, say, in the r band be consistent over the full survey; while absolute calibration requires that the r band magnitudes be convertible into physical units (Jansky). The relative calibration, again, is made easier by the multiple images of any given area of sky. The absolute calibration reduces to a single conversion factor per filter, but may require separate techniques to determine.

Astrometric calibration:

There is more experience with astrometrically calibrating large surveys over wide angles; it has been done both from space (Hipparcos) and the ground (USNO-B, the HST Guide Star Catalog, UCAC). The science goals in § 6.3.1 call for wide-angle astrometry good to a few milli-arcseconds, which again is helped tremendously by the repeat imaging and overlap of fields in any area of sky; there is a sense that any "reasonable" observing cadence will be adequate in this regard. Defining the set of standards to be used will be paramount, and has important implications on such issues as the saturation limit of the instrument. An early survey of the sky in multiple bands with the LSST will be crucial for this goal, as suggested in § 10.

Observatory location:

The extragalactic science goals of the LSST can be done from any site within latitude, say, $\pm 35^\circ$, to allow an appreciable solid angle coverage through the year. The Galactic structure and solar system goals put a premium on dual-hemisphere coverage. We did not discuss the relative merits of a Northern or Southern site for the LSST,

nor did we find an overwhelming case for simultaneous operation at two sites, one North and one South, thus allowing full coverage, e.g., of the ecliptic plane. It is worth pointing out in this context that the Pan-STARRS distributed aperture approach (Appendix C) does not require that all telescopes be at the same site (although the Pan-STARRS team currently plans to put all their telescopes at a common site).

More important than latitude is the seeing at the site. The weak lensing science goals of LSST require superb seeing, of order $0.7''$ median, and this must be one of the strong drivers in choosing a site for the LSST. While it further quantification is needed, the difference in science return between a site with median seeing of, say, $0.65''$ and $0.8''$ is likely to be large.

Of comparable importance is the weather on the site. The LSST is a photometric telescope, and the larger the fraction of photometric and clear nights, the more efficiently the survey will proceed. The fraction of perfectly photometric nights is not crucially important, given the large number of repeats of any given area of sky, of which only one exposure need be truly photometric; however, the fraction of reasonably clear nights (a concept that needs to be quantified!) will drive the rate at which LSST can gather data. The 8.4m LSST exposure time calculator (`http://www.ctio.noao.edu/lsst/etc/`) allows one to include a model for the weather, but this needs to be integrated with strategies for photometric calibration to determine what is possible at any given site. This is all made more difficult by the notorious $1/f$ nature of weather: there are very long timescale variations in weather conditions at any site that make historical weather records sometimes a poor predictor of the future. In addition, weather records at different sites often use different, and not properly quantified definitions of "clear" and "photometric", making the comparison of different sites problematic.

Exposure time for individual exposures:

This may seem to be a detailed issue associated with telescope cadence and observing strategy, but lies at the heart of the comparison, e.g., between the 8.4m LSST approach of a single large aperture (Appendix B) and the distributed aperture approach of Pan-STARRS (Appendix C). One needs to expose long enough that sky noise dominates over read noise; one clearly reaches this limit quicker on a large telescope than a small one, for a given pixel size on the sky and choice of filters. One is driven to shorter exposure times by considering the time over which asteroids, especially NEAs, are trailed. § 3 argues for exposure times as short as 10 seconds, which is adequate for all but the very closest asteroids. One needs the large aperture size of the 8.4m LSST to reach the sky limit in such a short exposure time for standard wide-band filters. Such a short exposure time also puts the onus on telescope structural design, as the settling time after a slew can become a serious limiting effect on the duty cycle. The Pan-STARRS discussion (Appendix C) points out that most NEA's will be seen at a distance of 1 AU or more, where it takes of order 30 seconds to move the width of a PSF. One can gain further by choosing a very broad-band filter in a focused survey for solar-system objects. As discussed in § 3, further work is needed to quantify the tradeoffs of exposure time, filter size, sky coverage, and completeness of the NEA survey.

Shorter exposure times also expand the discovery space for short timescale variability, although we did not identify a firm requirement on this front. In addition, for a given total exposure on a given area of sky, the shorter the individual exposures and the more of them there are, allowing one to average over, or determine the effects of, a number of the systematic effects that plague weak lensing measurements (see the discussion in § 7). Finally, we note that short exposure times puts the onus on telescope settling and camera readout times, which can add appreciably to the overheads.

Filters and Cadences:

Multi-color photometry is key to many, if not most of the science goals of the LSST. A fairly standard filter set, such as the SDSS $griz$, perhaps with a Y filter centered on one micron, which much improves the determination of photometric redshifts, satisfies most science needs discussed in this document. It is interesting that there is no strong need for a u filter; cf., the discussion in § 9.2. It has been suggested that the NEA search be done through a very broad filter, to increase throughput. However, it is quite difficult to do accurate photometry through a broad filter, and the refraction corrections to the astrometry are also difficult to calibrate; such data would not be useful for the other science goals. It has also been suggested that filters can be specially designed to be optimal for photometric redshifts following Budavári et al. (2001, AJ, 121, 3266); this needs further exploration.

Whatever the filter choice may be, calibration and cross-reference between different surveys will be much enhanced if the *same* filters are chosen between different large-scale missions. Pan-STARRS (Appendix C) will have surveyed the sky and defined a tight grid of photometric calibration standards by the time a full-scale LSST comes on line, and it makes sense to take advantage of this by matching the LSST filter choice to that of Pan-STARRS. Discussions between the 8.4m LSST and the Pan-STARRS team are underway to optimize filter choice for both surveys.

The Universal Cadence described in Appendix A is driven largely by the desire to follow NEA's in the confusing background of main-belt asteroids. We have seen in § 8.1 that it doesn't give fine enough wavelength or temporal coverage of supernovae to address the core supernova science goals. Further work is needed along these lines, including detailed end-to-end simulations of the scientific return for a given program given various cadences. There is a similar imperfect synergy between the ideal cadence for the deep Kuiper Belt Object search (§ 4.4) and the high-redshift supernova search (§ 8.3); again, the latter wants a four filters rather than two, and quite a bit faster cadence, albeit perhaps over fewer fields.

Data Systems and Data Distribution:

The data system of the LSST will be absolutely crucial to its success. With a data rate measured in terabytes per night, with synoptic science at the core of its goals, and with the need for continuous quality assurance to diagnose problems and to set observing strategy, it is self-evident that the LSST data will have to be processed in essentially real time (i.e., within 24 hours of it having been taken). The definition of 'processing' needs some real thought, but will have to include at least:

- Flattening and sky-subtracting the data.
- Determining the photometricity of the data (see above).
- Determining at least a preliminary photometric calibration.
- Astrometrically calibrating the data.
- Warping the image onto a standard projection.
- Differencing the image from a fiducial.
- Identifying from the difference image (and/or perhaps from a catalog of identified objects on the original frame) all variable and moving objects, and reconstructing their photometry and positions.
- Updating a database with this information.

- Deriving orbits for all solar-system objects found in the field.
- Identifying any noteworthy time-critical phenomenae (NEO's, supernovae, microlensing events, etc.).
- Using the data to assess the health of the telescope and instruments: Electronic problems in the camera? Collimation problems in the optics? Jitter in the telescope? Telescope throughput? Etc., etc.

All these data need to be made readily available to the entire scientific community. As we describe in § 11, these data are valuable to the extent that they are broadly available to the community, in a form that enables science to be done directly off the database.

On a longer timescale, of course, there is additional processing that must be done: refinement of the photometric and astrometric calibration grids, coaddition of the fields and analysis of the resulting images, and so on.

There is also a non-trivial amount of algorithm development that must take place: some issues that come to mind are correcting for differential chromatic aberration before co-adding or subtracting images, linking together orbits of asteroids in fields containing hundreds of moving objects, optimal coaddition of data, optimal photometry and shape measurements of faint sources, and so on.

The LSST data are quite complex, and making it available to the scientific community will be non-trivial. The data range from the unprocessed images directly off the telescope, to light curves of objects in multiple bands, to catalogs of image shapes measured off stacked images.

The LSST is defined by more than just the data that it will produce: its ultimate product is the science results. The LSST project must find ways to support the core science goals of the survey, rather than just making the data available and letting the "community" do the science. In addition, using the data to do science, in particular, the technically challenging science goals outlined in this document, is by far the best way to commission the system and perform quality assurance. The process of looking for fast-moving objects which are candidate NEO's will highlight the presence of image artifacts, and will push the orbit-fitting techniques to their limit. Weak lensing analyses will necessarily require an exquisite determination of the determination of the image Point-Spread Function, and will focus attention on systematics in the optics, camera, telescope, and image processing that limit its quality.

Comparing Missions, and the Scientific Landscape in 2012

The present document has focussed on the scientific capabilities of an LSST, but there are broader questions that we have only touched upon here, and which will become the focus of the LSST SWG now that the present document is complete. In particular, we need to understand the scientific landscape that will exist in ~ 8 years (a perhaps realistic time for the LSST to see first light). At that time, the 4-telescope Pan-STARRS will have been operating for several years, various supernova surveys getting started now will have discovered hundreds of objects, and so on. To what extent will the LSST goals described in the present document be completed by that time? What new questions and new approaches will become apparent in that time to modify the goals? These are major questions for the SWG to work on in the next years.

A closely related question is the optimal implementation of the LSST concept. We described in the introduction two possible missions, a single monolithic telescope of aperture 6.9 meters, and a series of smaller (1.8m) telescopes (Pan-STARRS, or "Multi-STARRS, an expansion to 15-20 telescopes). We have not addressed in this report which approach is superior. The appendices include summaries of each of these approaches. In Table 12 below, we compare the basic parameters of the 8.4m LSST, Pan-STARRS, and two other proposed missions of substantial étendue with overlapping science goals to those described in the present document: the Super-

Nova Acceleration Probe (SNAP; `http://snap.lbl.gov`), a space-based mission to discover Type Ia supernovae to redshift 1.7, and the Dark Energy Camera (DECam; `http://home.fnal.gov/~annis/astrophys/deCam`), a proposed wide-field imaging camera for the CTIO 4-m telescope.

There are a number of figures of merit one can derive for various scientific goals: for example, the sensitivity to point sources is proportional to étendue divided by the area of the PSF, with additional factors for background sky brightness; efficacy of a survey for variability studies will include factors depending on individual exposure time, and so on. The SWG plans to quantify these figures of merit in some detail in future work.

Table 3. Comparison of parameters of several proposed wide-field instruments

	8.4m LSST	Pan-STARRS	SNAP	DECam
# Telescopes	1	4	1	1
Effective Aperture (m)	6.9	1.44	2.0	3.57
Field of view (deg^2)	10	7.07	0.34+0.34	3
Pixel Size (arcsec)	0.2	0.3	0.10,0.18	0.27
Effective PSF (arcsec FWHM)	0.7	0.7	0.15 (at 700 nm)	1.0
Étendue (m^2deg^2)	360	46	2.1	30
# Filters	5	6	9	4
Exposure Time (seconds/exposure)	10	30-60	67,200,300	100
IR?	No	No	Yes	No
Spectroscopy?	No	No	Yes	No
Proposed First light	2012	2006 (first telescope)	4 years from start of funding	2008
Nominal Project Lifetime (years)	10	10	3	5

Note. — Numbers are drawn from the websites of the various projects, and in some cases (e.g., number of filters) are not yet set in stone. Note that SNAP, being in space, will have a ~ 7 times darker sky in r than the ground-based missions, and an even larger advantage at longer wavelengths. It will also be able to operate continuously, and thus will have a roughly three times larger duty cycle than 8.4m LSST and Pan-STARRS. Of the three ground-based projects, the 8.4m LSST and Pan-STARRS will be dedicated facilities, while DECam will be granted 30% of time on the CTIO 4-meter telescope over the project lifetime.

A. A Possible Universal Cadence

Strict optimization of each of the numerous science programs that LSST will enable would certainly result in the same number of observing strategies. Here we develop a universal cadence that would still allow one to address most programs in a nearly optimal way. The cadence described below assumes an étendue closer to that proposed for the 8.4m LSST ($\sim$265 m^2deg^2), but some of its philosophy may also be applicable to Pan-STARRS.

We start with the assumption that the most important goal of LSST is to detect solar system objects that may be on a collision course with Earth, and show that the main difficulty in finding such objects will be the background of main-belt asteroids. We demonstrate that there are intrinsic time scales set by the sky density and proper motions of main-belt asteroids that any cadence should address, and argue that multi-color photometry can be efficiently integrated in the proposed cadence. The resulting cadence has a number of desirable properties, and in particular, samples a wide range of time scales that are necessary for time domain science. Equally important, the proposed cadence is invariant to time translation and reversal, a feature that is desirable for a massive steady-state synoptic sky survey.

A.1. The Constraints on Time Interval Between Two Revisits

Based on the extrapolation of SDSS main-belt asteroid (MBA) counts from $r = 21.5$ to $V = 24$ asteroids have a peak sky surface density on the ecliptic of 226 deg^{-2}. This density is about two orders of magnitude higher than the expected density of potentially hazardous asteroids (PHA), and thus MBAs must be efficiently and robustly recognized in order to find PHAs.

The MBA sky surface density translates into a mean distance between two objects of 2.3 arcmin. This distance is much smaller than the typical MBA motion in 24 hours (3-18 arcmin), therefore requiring at least two observations during a single night. With a pair of observations closely spaced in time, the mean displacement is much smaller than the typical distance between two objects, and the recognition and linkage of moving objects becomes trivial. It is mathematically possible to link observations and constrain orbits even with a single detection per night, given a sufficient number of observations. However, due to a large number of expected detections, and six-dimensional orbital parameter space, this is a formidable task. With two observations per night, the velocity vector can be constrained sufficiently accurately that the night-to-night linkage becomes trivial. Such straightforward and fast processing method is of crucial importance for real time analysis.

What is the optimal time interval between two observations in a given night? The linkage confusion, C, increases with t, where t is the time interval between the two observations, as

$$C = 3.5\% \left(\frac{t}{15\mathrm{min}} \right)^2 \tag{A1}$$

This is the upper limit obtained using the sky surface density of 226 deg^{-2}.

Times shorter than 15 minutes result in smaller linkage confusion. A lower limit for t can be set by requiring a 5σ detection of a typical KBO (assuming an object at 100 AU moving at 22 mas/min), and is equal to $\sim$ 10 minutes. Thus, the above considerations results in a very narrow range of allowed t. Hereafter, we adopt t=15 min.

A.2. A Single Night Strategy

Assuming two individual exposures of τ sec (to reject cosmic rays), and s seconds for slewing (readout occurs during slew), the number of fields of view (FOV) that can be observed within t minutes, i.e. until the first FOV

has to be revisited, is

$$N_{FOV} = \frac{60\ t}{2\tau + s} \tag{A2}$$

With benchmark values of $\tau = 10$ and $s = 5$, we find $N_{FOV} = 36$ (for $s = 8$, $N_{FOV} = 32$). Note that the adopted s represents an "effective" value because the slewing time in the azimuth direction may be about a factor of 2 longer than in the altitude direction (due to frequent dome movement).

The only hardware-dependent assumption so far is that $V \sim 24$ can be reached with the integration time of 2τ sec. The second hardware-dependent assumption employed here is that the FOV has a diameter, D_{FOV}, of 3 deg (area of 7 deg^2). With such an FOV, a possible strategy is to observe $N_{FOV}/2 = 18$ FOVs along a line of constant ecliptic longitude[9], with a step of $\sqrt{3}\,D_{FOV}/2$ deg. Then the next $N_{FOV}/2$ FOVs are observed in the opposite direction, with a longitude offset of $3D_{FOV}/4$ (see the top two panels in Figure 22). After all N_{FOV} FOVs are observed, the same fields are reobserved, and then a new set of N_{FOV} FOVs are observed for the first time (bottom left panel in Figure 22). Assuming a night T_h hours long, the total number of observed FOVs (two visits each) is

$$N_{obs}^{FOV} = \frac{30\,T_h\,N_{FOV}}{t} = 72\,T_h \tag{A3}$$

Adopting $T_h = 9$, we find $N_{obs}^{FOV} = 648$. With the effective fraction of the observed area of $3\sqrt{3}/(2\pi) = 0.83$, the total observed area is

$$A_{obs} = 0.83\,N_{obs}^{FOV}\,A_{FOV} = 59.8\,A_{FOV}\,T_h \tag{A4}$$

With the adopted values of t, τ, s, and D_{FOV}, the sky is observed at a rate of 418 deg^2 hour^{-1}, with two visits per FOV separated by t minutes. With $T_h = 9$, the total unique area observed in a night is 3765 deg^2 (total observed area is 4536 deg^2). Thus, the same area on the sky could be comfortably observed every three nights. Work remains to be done to confirm that this is adequate for robust orbit determination, given the possibility of bad weather, bright time, etc?

A.3. Night-to-Night Strategy

Two observations obtained t minutes apart in a single night, constrain the position and velocity vector of a moving object. How long into the future can the position of such an object be predicted before the confusion with other moving objects becomes intolerable?

The uncertainty in predicted position, employing a linear extrapolation along the velocity vector, is

$$\Delta = 5\,\mathrm{arcsec}\left(\frac{\sigma_A}{50\,\mathrm{mas}}\right)\left(\frac{15\,\mathrm{min}}{t}\right)\left(\frac{T_{nn}}{1\,\mathrm{day}}\right) \tag{A5}$$

where T_{nn} is the time in days between the two repeated observations of the same field (night-to-night), and σ_A is the astrometric accuracy ($\sim$ 50 mas, smaller values strengthen the following argument). Even if T_{nn} is as long as 30 days, Δ is as small as 2.5 arcmin, so the velocity measurement error is not the limiting factor – the limit is set by the non-linearity of asteroid orbits.

The practical limit on the maximum acceptable value of T_{nn} depends on how sophisticated is the orbit prediction method based on two positions obtained 15 min apart. Using a simple circular orbit approximation, this limit is about 5-7 days. With a more sophisticated treatment, such as a Bayesian approach developed by

[9]Due to different slewing times in the azimuth and altitude directions, it is likely that the tiling strategy discussed here can be further optimized.

Bowell and collaborators (private communication), this limit could be extended to several weeks. Thus, it is likely that a minimum of two nights per field and per dark run should suffice.

The main conclusion is that repeating the observations of the whole sky every three nights, with a preference for reobserving after a stretch of bad weather given to the Ecliptic plane area, should allow a robust and efficient recognition and linkage of moving objects.

A.4. Multi-band Strategy

In the preceding sections it was assumed that each field would be visited at least twice in a given night. These observations could always be done using the same filter, or no filter at all. However, it is advantageous for practically all science programs, including solar system science, to obtain multi-band photometry. This can be achieved with the proposed cadence by switching filters between the two observations obtained in a single night. One possibility is to utilize pairs (r, g), (r, i), (r, z) for, say, three visits during a dark run. This scheme helps to avoid the effects of rotation on asteroid color (note that averaging multiple random observations would take 10-20 epochs per band to get the same asteroid color accuracy as two observations obtained 15 min apart), and produces deeper data in one selected band (here r). Such deeper data greatly improve the faint source detection.

One case where the faint source detection is of crucial importance is the use of optical catalogs obtained by LSST to identify sources detected at other wavelengths (e.g. X rays, infrared, radio). For example, only a third of radio sources detected by the FIRST survey have optical counterparts detected by SDSS (Ivezić et al. 2001b, 2002c). Deep optical surveys of small areas indicate that practically all FIRST sources should have $V < 26$. Thus, using LSST and FIRST, it will be possible to construct a catalog of over a 1 million sources with both optical and radio detections!

A.5. The Sampling of Different Time Scales

Due to the overlap between FOVs, 17% of the observed area represent multiple observations with a variety of time scales. With the choice of free parameters used here, about 5% of the area ($\sim$200 deg^2 per night) would be reobserved with the time interval of 25 sec. Another 10% of the area would be reobserved with a fairly uniform sampling of time scales ranging from 25 sec. to 15 min. This range of time scales is practically terra incognita with the currently available data, and it may uncover and provide robust statistical measurements for many new transient phenomena (§ 5). For example, the Deep Lens Survey (§ 7) is finding transients that last only 10–30 minutes, appear unresolved, and have no precursor object.

Of course, there is a lot of freedom in fine tuning these time scales (and also to go deeper than $V \sim 24$) by changing the various free parameters adopted in this section.

The cadence described in this section is clearly only a first step. We have seen throughout this document that this cadence is roughly what many science programs require; this all needs to be properly quantified and simulated, using, e.g., the 8.4m LSST operations simulator developed by C. Smith and K. Olsen. There are notable exceptions, such as the deep KBO survey described in § 4 and the deep supernova survey in § 8, which this cadence does not address.

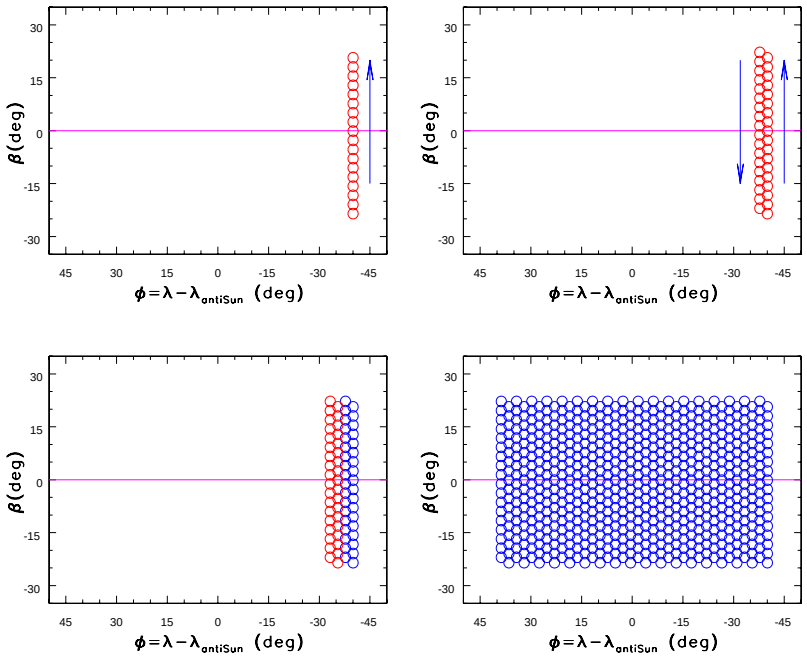

Fig. 22.— The four basic steps in an universal cadence strategy. The sky is observed along a line of constant ecliptic longitude, moving towards the north ecliptic pole with a step of $\sqrt{3}\,D_{FOV}/2$, where D_{FOV} is the field-of-view diameter (top left panel). After $N_{FOV}/2$ fields are observed, where N_{FOV} is the number of fields that can be observed until the first field has to be revisited, the next $N_{FOV}/2$ fields are observed in the opposite direction (top right), with a longitude offset of $3D_{FOV}/4$. In the next step the same N_{FOV} fields are reobserved, and then the next set of fields is observed for the first time (bottom left). With reasonable assumptions for various free parameters, the total observed area in one night is shown in the bottom right panel.

B. The 8.4m LSST System Reference Design

This Appendix was written by J. Anthony Tyson, member of the LSST SWG and Director of the 8.4m LSST

A baseline design for the 8.4m LSST has been developed which will allow it to achieve many scientific goals simultaneously with high quality data in a single new technology optical system. The project has raised sufficient support to proceed with engineering for all systems, including software. A nonprofit corporation manages the fund raising and the project. The data will all be public, without any proprietary time. Data access solutions will be developed, and there will be several data products. A hierarchy of photometric catalogs, updated transient and moving object databases, and the full imaging database will be available. In addition, at least two science data products will be produced: an optical burster and orbit database, and weak lens shear maps vs redshift. The project timeline calls for first light in 2011 and operations in 2012.

This chapter assumes a usable field of view of 3°; however, the 8.4m LSST collaboration is now planning a larger field of view, of 3.5°, or 10deg^2. This gives an ètendue of 360 m^{2}deg^2. Note that this updated value is reflected in Table 12 in § 12.

B.1. The 8.4m Telescope

B.1.1. Optical Design

The optical design for the 8.4m LSST (Figure 23) is based on a concept by R. P. Angel et al. (2000) which modifies the Paul-Baker three-mirror telescope to work at large apertures. Seppala (2002) further developed the Angel design, simplifying the aspheric surfaces and achieving a flat focal plane. The current design employs three aspheric mirrors (an 8.4-m diameter primary, 3.5-m secondary and 5-m tertiary) which feed a three-element correcting camera resulting in a 3° ($\sim$ 7 deg^2) circular field of view (FOV) covering a 55-cm diameter flat focal plane. This design has an étendue of 266 m^{2}deg^2, satisfying the requirements of a system which will survey over 14,000 square degrees of sky in several filters multiple times per lunation. Particular attention has been paid to image quality and stability required for weak lensing. These optics deliver a PSF with Full Width at Half Maximum (FWHM) better than 0.2 arcsec over the entire 7 square degrees in all bands from 350 nm to 1000 nm.

The primary and tertiary mirrors are both concave and are similar in diameter and f/number to mirrors already produced by spin-casting. The large convex secondary will be a structured, light-weighted mirror made from a low expansion glass (e.g. Zerodur). A high-performance, multi-layer coating designed at Lawrence Livermore National Laboratory will be applied to each mirror to maximize integrated throughput. The three refractive lens elements will be made from optical-grade fused silica. The convex surfaces of the first two lenses, L1 and L2, are eighth order aspheres while all other surfaces are spherical. The filter is a zero-power meniscus which keeps the chief ray normal to the surface everywhere across the field of view; this ensures a uniform band-pass across the field from multi-layer dielectric coatings – important for precision photometry.

B.1.2. Telescope Design

It must be possible to point the telescope quickly ($<$ 5 seconds) and repeatedly to adjacent field locations. Because of its compactness and design maturity, an Alt-Az mount configuration is a logical choice; two conceptual designs being considered are shown in Figure 24.

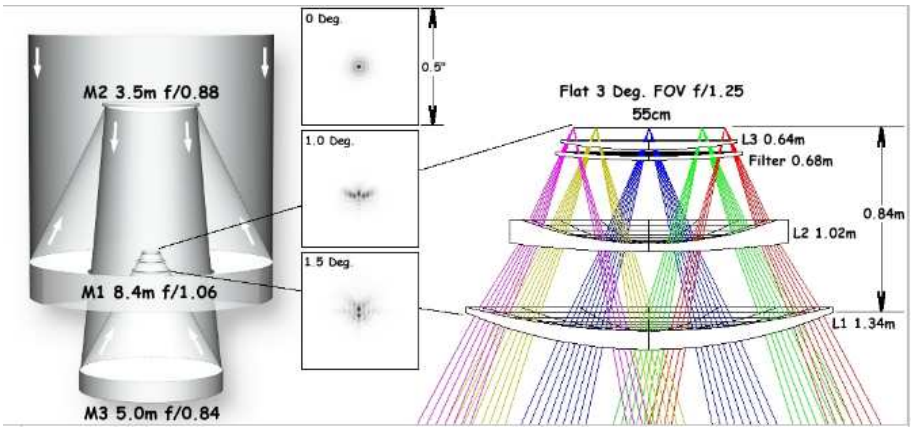

Fig. 23.— The overall optical path for the 8.4m LSST optical design (left) along with a more detailed view of the corrective camera optics (right). Insets show the ray-tracing response to a point source at 650nm for three field radii.

Table 4. Optical Design Summary

Optical Configuration:	3-mirror modified Paul-Baker
Aperture:	8.4 m
FOV:	7.1 deg^2
étendue (AΩ):	266 $m^2 deg^2$
Wavelength Coverage:	300 – 1100 nm
Image Quality (80% EE diam.):	$< 0.25''$ (BVRI), $< 0.35''$ (U) FWHM
Effective Clear Aperture:	7.078 m (6.9 m incl. obscuration)
Final f-Ratio:	f/1.25
Plate Scale:	50.9 microns/arcsec

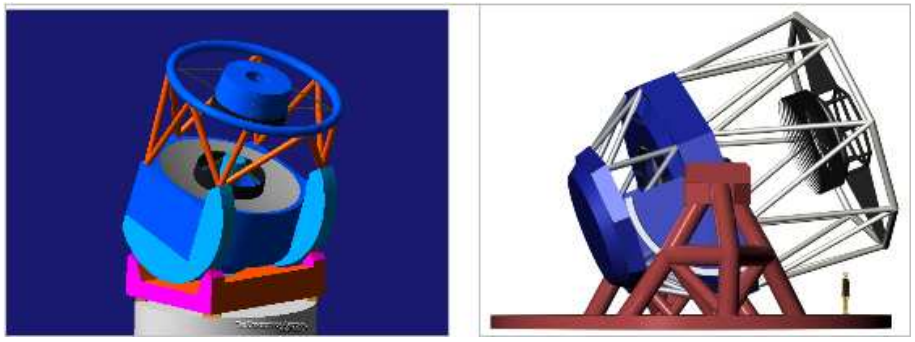

Fig. 24.— 8.4m LSST Alt-Az Telescope concepts. At left is a dual "C" ring concept similar to the LBT and MMT, by Davison (UA, Steward Observatory). The right panel shows a concept based on a more-traditional fork configuration by Claver and Muller (NOAO).

B.2. the camera and focal plane assembly

B.2.1. Camera Design

The 8.4m LSST camera is a wide-field optical (0.35-1 μm) imager designed to provide a 3° FOV with better than 0.2 arcsecond sampling. The image surface is flat with a diameter of approximately 55 cm. The detector format will be a circular mosaic providing approximately 2.3 Gigapixels per image. The camera includes a filter mechanism and, if necessary, shuttering capability. It is positioned in the middle of the telescope where cross-section area is constrained by optical vignetting and heat dissipation must be controlled to limit thermal gradients in the optical beam. The camera must produce data of extremely high quality with minimal downtime and maintenance.

The camera concept currently under development is shown in Figure 25. The design shows a "dewar within a dewar" structure. The inner dewar contains the detector array, held at a temperature of -40 C in order to achieve desired detector performance. The refractive element L3 (Figure 23) serves as the window of the inner dewar, while L1 serves as the window for the outer dewar. The outer dewar houses L2, the filters, and the filter exchange mechanism, which can accommodate four 60-cm filters. This mechanism uses a novel approach to adapt to the extremely tight space constraints. The camera mechanical mount will provide proper support and registration to the telescope and incorporate provisions to actively adjust the camera position and orientation to compensate for alignment variations with telescope elevation. In addition, the camera axial position must be adjustable to optimize focus at different filter wavelengths (the axial position of L2 must be similarly adjustable).

B.2.2. Detector Arrays

Recent advances in CMOS and CCD + ASIC hybrid imagers can be applied to this 2.3 gigapixel camera. An array built from a mosaic of 2K modules is preferred. Parallel multiplexing many discrete modules allows for fast readout, which will be critical for efficiency given the required short exposure times. The clocking electronics will be integrated with the individual detectors, and there are several attractive options for analog and digital ULSI packaging that minimize the number of interconnections. Each module will consist of a thick silicon detector for high QE over the full wavelength range from 350 nm to 1050 nm. The 8.4m LSST's large focal plane and short exposure times make the traditional approach of CCD plus mechanical shutter difficult to implement. LLNL has developed a shutter concept for the 8.4m LSST that would be capable of at least 10^6 exposures, or more than one year of 8.4m LSST operations. However, hybrid CMOS detector arrays with integrated ASIC electronics, originally developed for IR arrays, are now being produced for visible-wavelength applications and would eliminate the need for a mechanical shutter. Work is in progress on both monolithic and hybrid CCD array + readout electronics solutions to our module requirements. A 8.4m LSST exposure time calculator has been developed, based on the expected performance of these detectors and the telescope optics: `http://www.ctio.noao.edu/lsst/etc/`.

B.3. 8.4m LSST Data System

Software is arguably one of the most challenging aspects of the 8.4m LSST. Its data management system must process and store more than 6 Tera-pixels per night, roughly the same as the whole 2MASS survey. Even with current computing and storage technology, it is possible to handle 8.4m LSST data volumes and rates, though not yet with the full real time response we will require. In complexity as well, 8.4m LSST exceeds previous astronomical surveys but by a manageable factor. Most of the 8.4m LSST's new complexity derives from the extraordinarily broad range of science it will make possible. For the first time, astronomers determining NEO

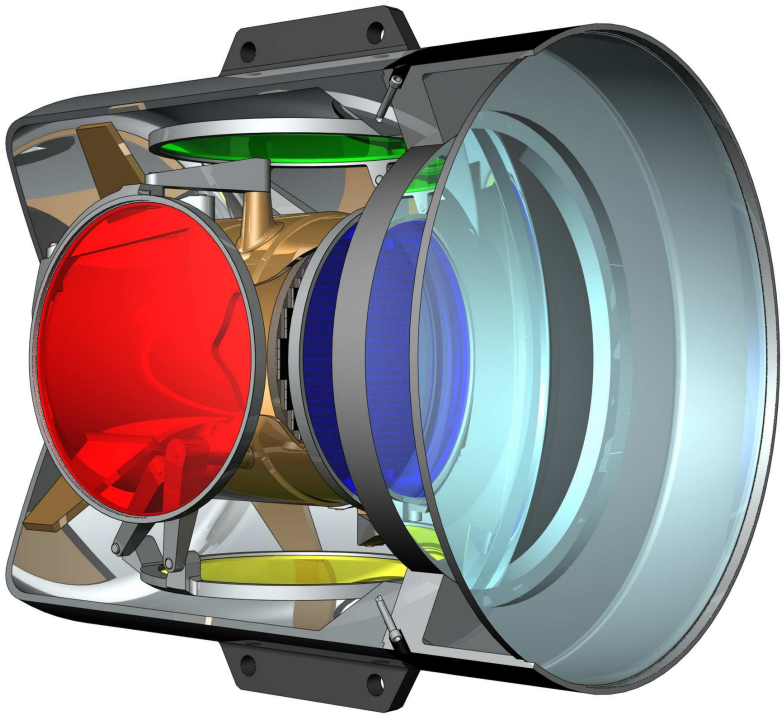

Fig. 25.— A sectional view of the camera dewar-in-dewar concept. The outer dewar houses the refractive elements for the wide-field correction, as well as the filter mechanism. The inner dewar holds the focal surface with detectors and interface electronics.

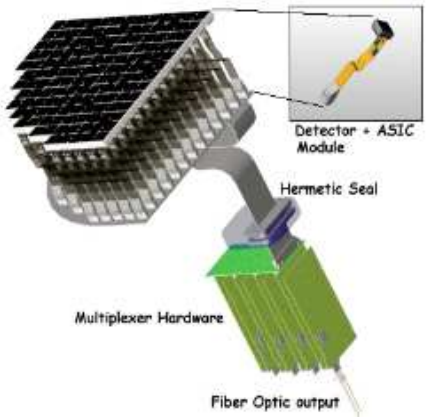

Fig. 26.— A quarter section of the detector array concept.

orbits will be sharing data structures with those determining weak lensing parameters and those classifying variable stars; this will inevitably complicate data structures. In what follows, we sketch a possible 8.4m LSST data system architecture, not as a design proposal, but as a basis for discussing the software issues that we face.

B.3.1. 8.4m LSST Data-Flow

Based on the scientific drivers and experience gained from recent large surveys, the baseline data flow diagram is shown in Figure 27. The data management system must move more than 2 Gigapixels of data from the camera into the data analysis system in 1-2 seconds: a rate of up to 5 gigabytes per second. These data must then move through pipelines that remove instrumental signatures, then through analysis pipelines, and into long term storage in about 2 seconds, with a hard upper limit of 10 seconds. The system will also put the results of analysis into various databases, make quality assurance data available to the telescope control system in real time, generate prioritized lists of transient astrophysical phenomena by comparison with previous data, and make this information available to the telescope scheduler and the public in near real time with a latency of less than 30 seconds. The flow shown here does not include Data Mining and User Portal activities.

B.3.2. 8.4m LSST Software Architecture

The 8.4m LSST data flow suggests a four-tier software architecture:

- A Telescope Control System (TCS), which manages the telescope and the camera system.
- A Core Data System (CDS), which processes and stores the camera pixel data stream, provides the infrastructure required by algorithms that access images, and implements the databases that store object data.
- An Image Analysis System (IAS), which further processes basic camera images, derives object information from those images, and detects transient events.
- A Science User System (SUS) that enables scientists and the public to use the archives and catalogs generated by the IAS and (under strictly controlled circumstances) allows access to the facilities of the CDS for further pixel-level processing.

B.4. Operations Simulations: A 8.4m LSST Design Reference Mission

The high étendue of the 8.4m LSST delivers sufficient pace on the sky to enable simultaneous achievement of multiple science goals. Exposures will be ten seconds, and will reach sky limit at 24 AB mag (10 σ). How well can this work in practice? A "design reference mission" in which all the properties of the facility (telescope, optics, camera, filters, observing algorithm) and site characteristics (seeing, weather, moon) are fully simulated will be necessary. A preliminary DRM using a five-band filter set (*bgriz*) covering the wavelength range 0.4–1 micron shows the following:

- An 8.4m LSST with étendue over 250 is required for the execution of the science programs.
- The 8.4m LSST in a good site can achieve these goals.

The simulation had as its goal to find an observation sequence achievable in ten years or less that is optimal for near-Earth asteroid (NEA) searches, while accumulating over 200 ten second exposures in each of five bands over 15,000 square degrees or more. The sky coverage should be tiled in a way that provides the needed overlap between fields for science goals and for astrometric and photometric tie-down. Stacking the exposures gives a 10σ magnitude limit for point sources of roughly 26.7 (AB), and a surface brightness limit of 29 magnitudes in one square arcsecond.

We assume the universal cadence defined in Appendix A, and define a "NEA sequence" as three pairs of visits of a given point in the sky separated by 2 and 6 days within a single lunation (defined to be the period from one full moon to the next). A pair, following Appendix A, consists of visits spaced by ~ 15 minutes, one with two back-to-back exposures in r, and the other with two back-to-back in b, g, or (rarely) i, chosen in such a manner to give similar total exposure times in all bands over time. The following are the further assumptions which went into this simulation.

1) Do not observe for ~ 3 nights at each full moon. 2) We take the site to be Cerro Pachon in Chile (because the necessary weather data were readily available), with seasonal weather pattern for that site (i.e. we use actual month-by-month clear night statistics in a random weather model). 3) Observe only in i or z bands when the moon is up and phase is larger than quarter (bright sky). 4) Observe in b only when the moon is down (darkest sky). 5) Observe in g or r either when the moon is down, or when the moon phase is smaller than a quarter (gray sky). 6) Maximum permissible air mass for observing is 2.0. 7) Observe only fields whose declination is less than 50° from the latitude of the observatory—all fields carry the same "prior" weight. 8) Observe only in "clear" conditions. 9) Do not count observations made in bright sky towards a NEA sequence. 10) If an NEA sequence is completed in any lunation, the field is still visited once every ~ 5 days—if dark, observe in b or g or r, if gray in g or r, and if bright, in i or z. 11) Each visit consists of two back-to-back 10-second exposures. 12) Penalties (telescope not exposing) are applied for large slews and for filter changes—no observations are made within 5° of the zenith—observations at lower air mass are favored when a choice exists. 13) Observe down to 10° Galactic latitude at zero longitude, going to $b = 0°$ at 180° longitude.

Extrapolating to ten years, and making the correction for uniform 15% field overlaps, we obtain 330 useful ten-second exposures in r, 332 in b, 316 in g, and 211 in each of i and z in every patch in 17,700 deg^2 of the sky. Using the 8.4m LSST exposure time calculator, we thus have 10σ limiting magnitudes for point sources (in 0.7 arcsec median seeing and mean air mass of 1.2) of $b = 26.8, g = 27.0, r = 27.0, i = 26.3$, and $z = 25.6$ AB mag. Figure 28 shows the achieved sky coverage and depth in all five bands. Thus with the étendue of the 8.4m LSST, the goal of achieving at least 200 ten-second exposures in each filter for each field in ten years is achieved.

A smaller part of the sky is covered if we require better than 70th percentile seeing in any of g, r, i, as well as avoidance of $\pm 20°$ of the Galactic plane at low longitudes: 14,000 deg^2 is covered to a depth of $g = 26.7, r = 26.8$, and $i = 26.2$ AB mag in that smaller dataset. Some depth in the excellent seeing data can be traded for more coverage. Thus it is possible to reach 15,000 deg^2 of high quality imaging in the critical bands for weak-lensing shear (§ 7) together with deep five-band photometry for photometric redshifts.

Note added in proof: The discussion of this chapter assumes a field of view of 3°; however, it now seems possible to increase this to 3.5°, increasing the ètendue to 370 m^{2}deg^2. Figure 29 shows the delivered image quality of the current design from ray-tracing of the optics, extending to a radius of 1.75°; it shows no serious degradation even to the edge of the field.

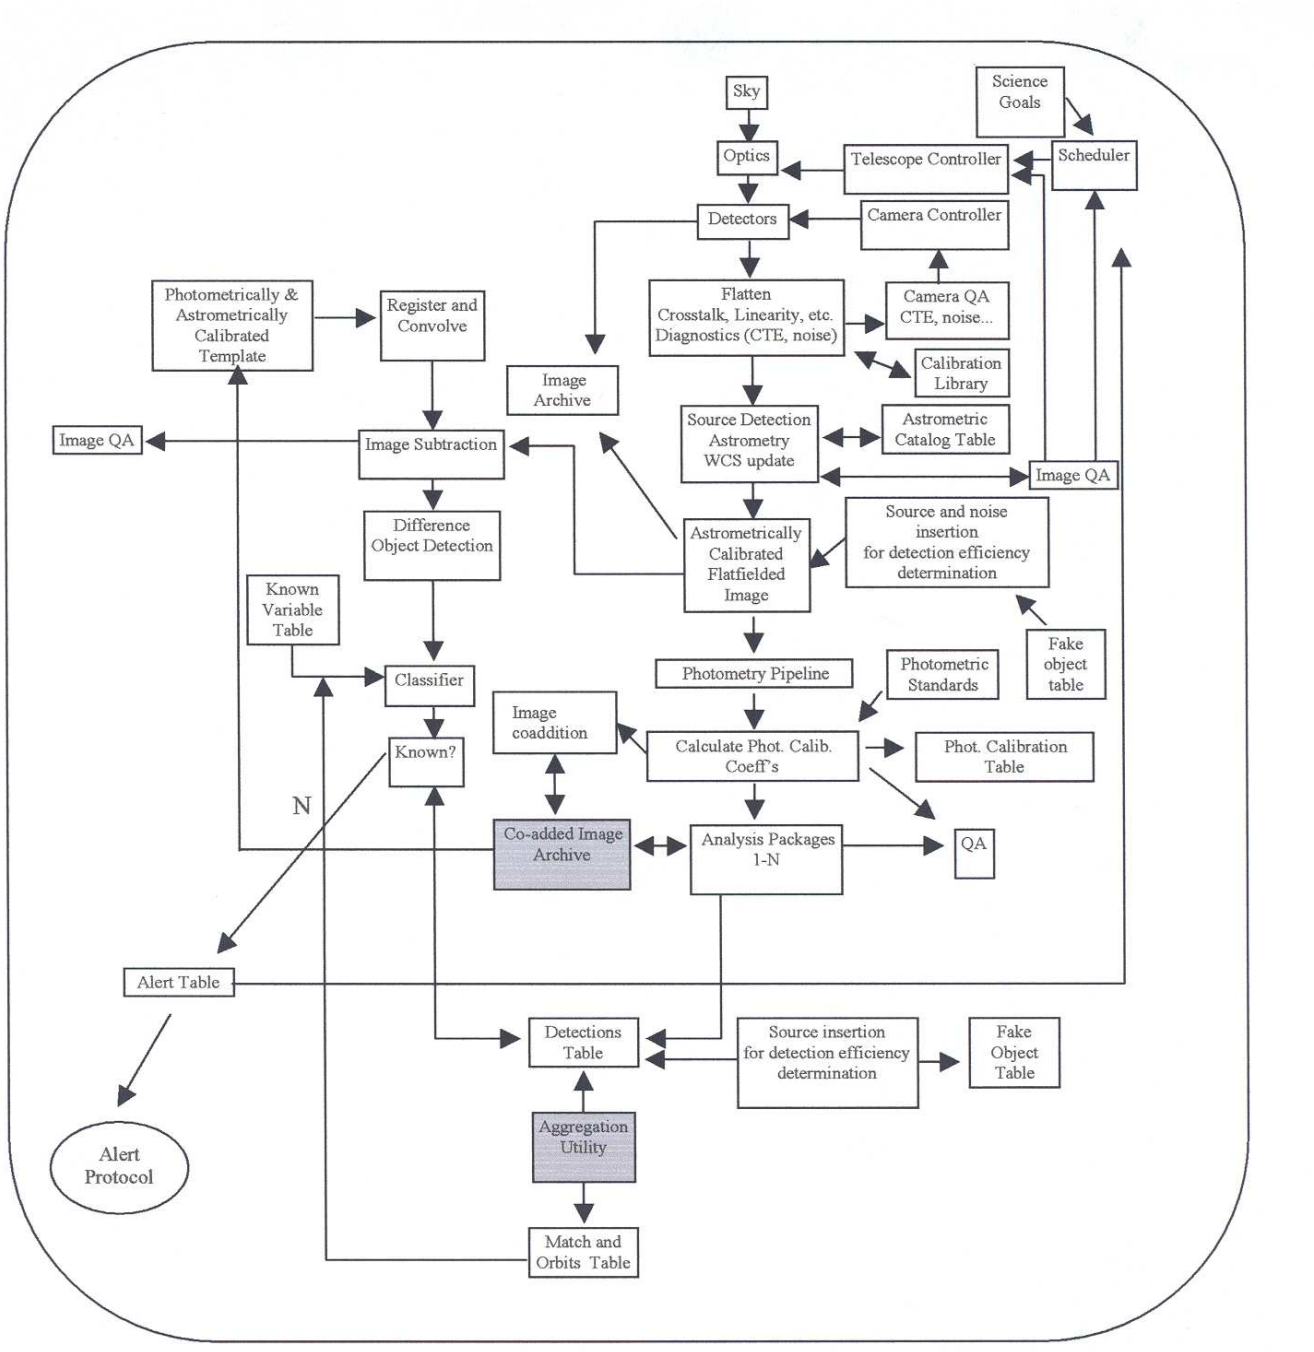

Fig. 27.— A model for the data flow in the 8.4m LSST.

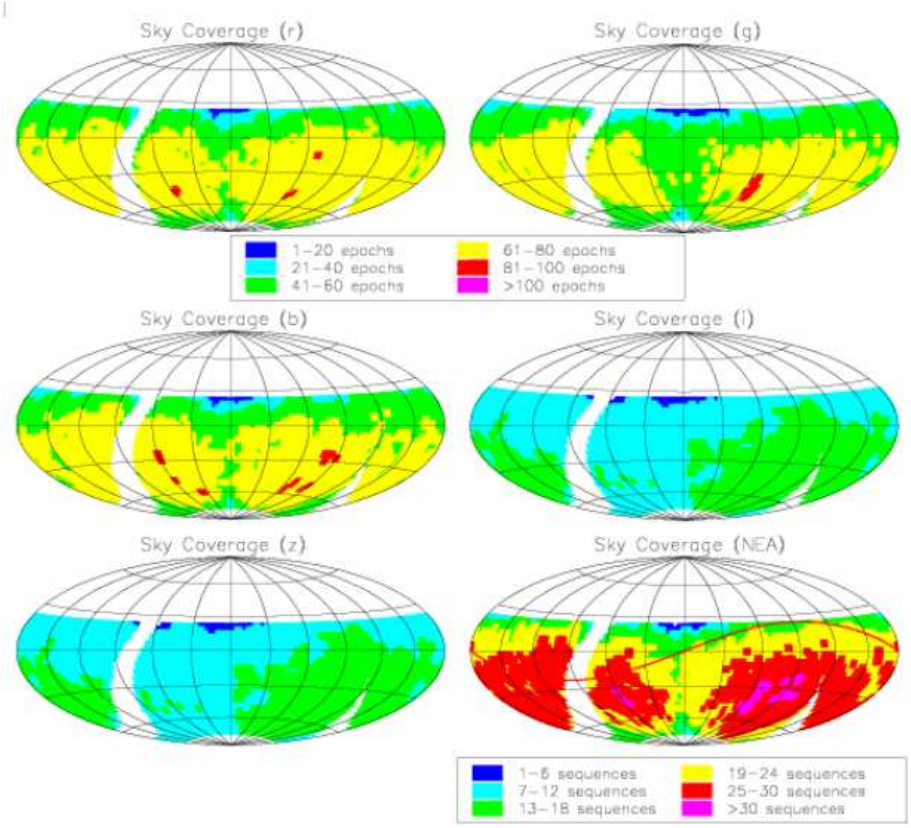

Fig. 28.— Exposure depth in *bgriz*, color-coded by cumulative number of exposures in four years and plotted in equatorial coordinates. The blank curved strip is the avoided area of the Galactic plane. For the 8.4m LSST this just reaches our goal of 26.5 AB mag depth over 19,000 square degrees in the co-add stacks in ten years, as well as rapid completeness for NEAs and good efficiency for transient detection. In these plots for the four year simulation, an "epoch" is one visit to a 7 square degree patch with a pair of 10 second exposures. The sixth plot (lower right) shows the NEA survey completeness over 4 years in terms of sequences completed. The red line is the ecliptic. A sequence consists of three pairs of visits separated by 15 minutes, 2 days, and 6 days within a lunation.

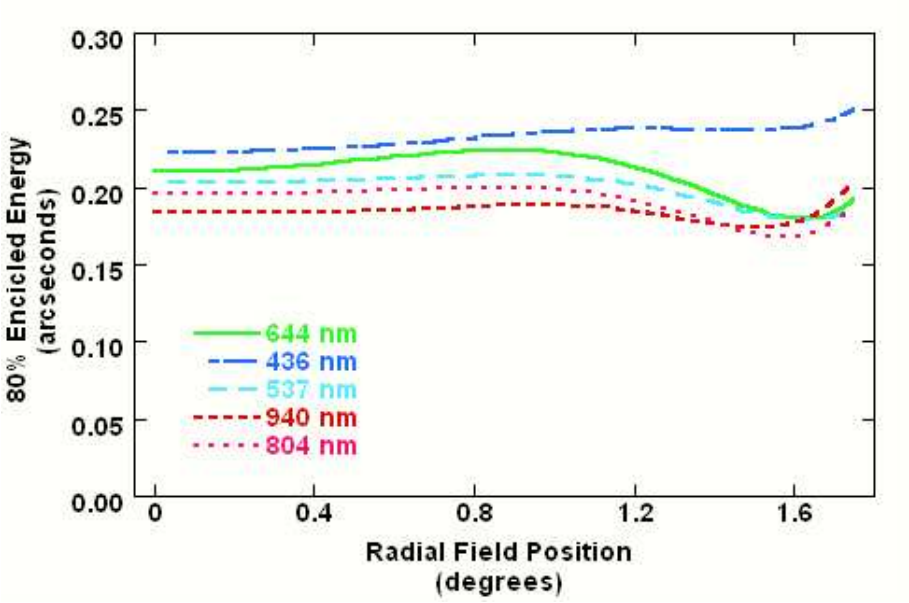

Fig. 29.— The radius enclosing 80% of the energy, as a function of wavelength and field position, due to the current (3.5°) design for the 8.4m LSST. Note the absence of serious degradation of the delivered image quality, even to the edge of the field.

C. Pan-STARRS — Overview and Project Status

This Appendix was written by Nick Kaiser, member of the LSST SWG and Project Scientist of the Pan-STARRS Project

The Institute for Astronomy at the University of Hawaii is developing a large optical synoptic survey telescope system; the Panoramic Survey Telescope and Rapid Response System, to be deployed on either Haleakala or Mauna Kea. Pan-STARRS will consist of an array of four 1.8-meter telescopes with very large (7 square degree) field of view, giving it an étendue of 46 $m^2\,deg^2$. Each telescope will be equipped with a 1 billion pixel CCD camera with low noise and rapid read-out, and the data will be reduced in near real time to produce both cumulative static sky and difference images, from which transient, moving and variable objects can be detected. Pan-STARRS will be able to survey up to $\simeq 6,000\,deg^2$ per night to a detection limit of approximately 24th magnitude.

Here we describe the Pan-STARRS system design, including the telescopes and detector designs, and the various survey modes needed to support the science goals. The data pipeline and the basic data products are reviewed. A detailed comparison is made between the performance metric for Pan-STARRS and the proposed 8.4m LSST design. We conclude with a discussion of the Pan-STARRS timeline and current status.

C.1. Pan-STARRS System Design

C.1.1. Telescopes

The observatory will consist of four 1.8m Cassegrain telescopes with three element wide-field corrector providing very good images over a 7 deg^2 field of view. The ZEMAX design meets the rather strict requirement of less than 10% image degradation (over good seeing conditions) from telescope aberrations. While not part of the preliminary design, an atmospheric dispersion compensator is now being considered. The focal length is 8 meters, yielding a plate scale of 38.5 microns per arcsecond. The layout of the optics is shown in Figure 30 and geometric optics spot diagrams are shown in Figure 31.

The design is fully baffled by means of an exterior baffle and a conical baffle in the converging beam, much as in the SDSS telescope design (cf., `http://astro.princeton.edu/PBOOK/camera/camera.htm`).

It is planned to use six filters; four closely modeled on the SDSS g, r, i, z, a Y-band filter centered around one micron, and a very wide 'solar system' filter optimized for detection of neutral color objects.

Low order quasi-static astigmatism (defocus, decollimation etc) will be monitored from pre- and post-focus images at the edge of the focal plane and will be actively controlled. Image quality considerations require that the detector surface remain flat to a precision of approximately $20\,\mu$m (this gives a image degradation about an order of magnitude less than that arising from finite detector resolution).

For individual telescopes of this size, slewing (over 3 degrees) and settling times are on the order of seconds. The telescopes will be actively guided at up to several Hertz using signals derived from guide stars distributed across the focal plane.

C.1.2. Detectors

The first batch of test devices currently being fabricated by MIT include 10 and 12 micron pixel sizes, as well as various options for gate layout for orthogonal transfer charge tracking. Both three-phase and four-phase (orthogonal transfer) devices are being tested. The final choice of detector design will be based on yield *vs.* performance considerations.

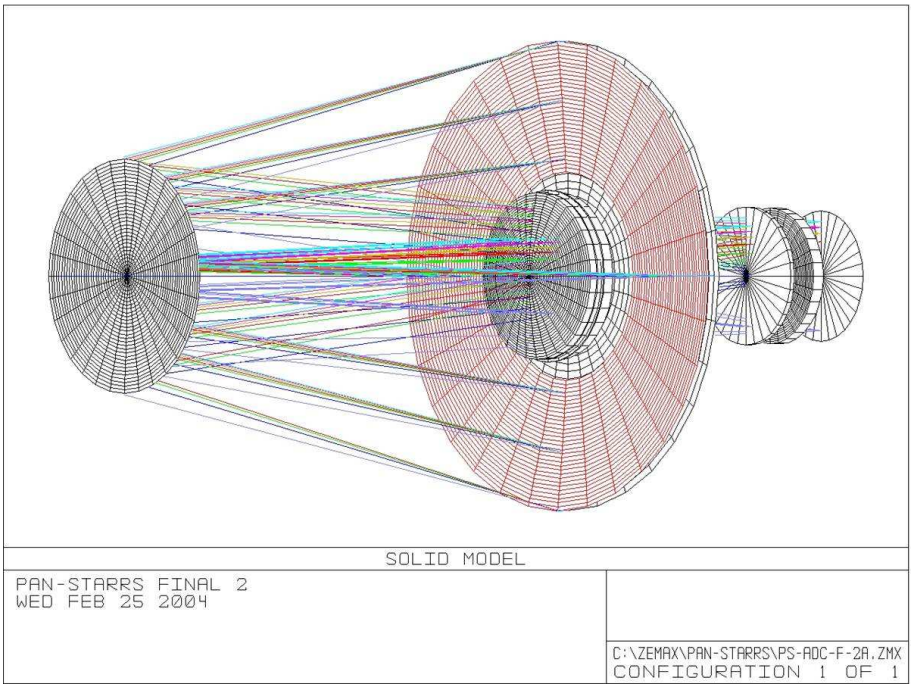

Fig. 30.— The Pan-STARRS optical design is a Cassegrain with a three element wide field corrector. One possibility being actively considered is to incorporate an atmospheric dispersion compensator in the first element of the corrector. The focal plane is flat.

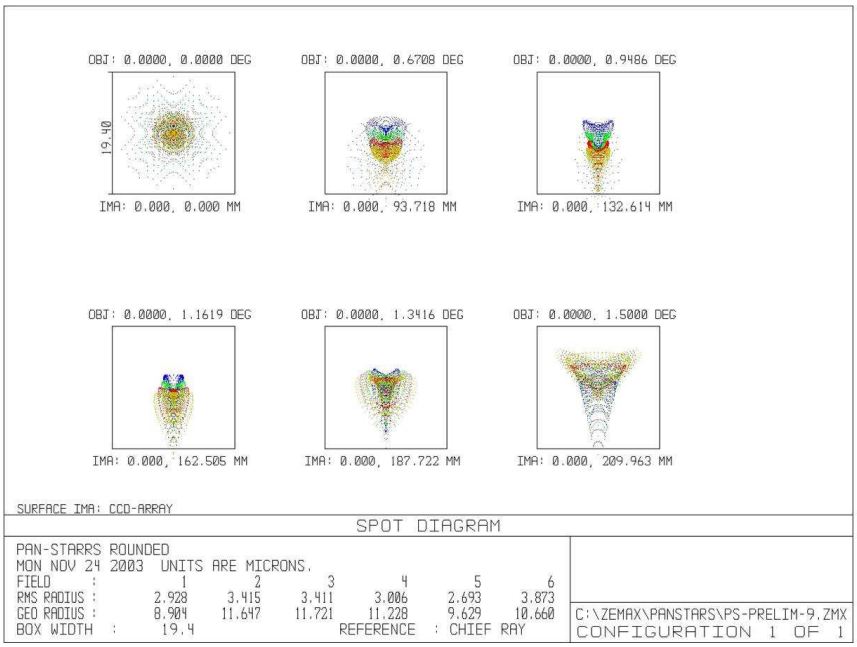

Fig. 31.— Geometric spot diagrams at various locations across the focal plane. The boxes are $0''.5$ on a side, so the image quality is very good indeed, and is essentially diffraction limited at wavelengths of interest.

A charge diffusion scale no greater than 5 microns is required. This produces an image degradation approximately three times that arising from the finite pixel size (for 12 micron pixels), but this is still a relatively minor degradation of the overall PSF even under good seeing conditions (see below). The required resolution has been demonstrated in the lab for thinned devices (40 micron thickness).

The design is an array of arrays: a focal plane consists of 60 devices (an 8 by 8 grid with the corner cells discarded) each approximately 5 cm on a side. A detector consists of an 8×8 grid of individually addressable cells each of approximately 512×512 pixels (this is for 12 micron pixels). The array detector design will allow greater effective yield, as an isolated defect will only disable one cell, and we can well afford to lose a small fraction of cells. The array detector design confers some other advantages; these include limiting the 'bleeding' from bright stars and the ability to rapidly obtain guide star centroids at any location on the focal plane without recourse to a pick-off mirror. The four-phase devices also allow independent and rapid on-chip fast guiding to improve image quality, as well as techniques such as 'PSF shaping' to do photometry on very bright stars.

Read noise of a few electrons rms with net read times of a few seconds is required. COTS analog-digital conversion devices have been shown to give acceptable noise ($\simeq$ 2.1ADU rms) and linearity ($\sim 1\%$) over the full dynamic range.

Thick test-batch devices are due around April 2004 and should provide useful yield estimates. Thinned devices will be available for full testing and characterization in the summer and a final decision on devices for the second and third lots to provide a full focal plane for the first Pan-STARRS telescope (PS1) will be made. Fine tuning of the design will involve a compromise between image quality and near IR sensitivity.

C.1.3. Fast Guiding

For small telescopes, a substantial increase in image quality can be obtained by fast guiding. Pan-STARRS will monitor positions of guide stars on the focal plane. A low frequency (~ 1 Hz) common mode motion signal will be used to guide the telescope, and the possibility of taking out somewhat higher frequency motions by control of the secondary mirror is being considered. For median Mauna Kea wind speed conditions (5 m/s) 90% of the image motion power arises from temporal frequencies below about 1.5 Hz. High frequency motions may be removed using on-chip fast guiding, where the shift for each cell is obtained by averaging over a collection of nearby guide stars, allowing for correction of motions that vary across the focal plane. This allows partial compensation of ground level seeing arising from heights on the order of tens of meters.

C.1.4. Surveys and Operation Modes

Pan-STARRS will perform several surveys. These include a 3π survey in five passbands ($grizY$), plus a selection of medium-deep and ultra-deep survey fields, also to be observed in five passbands[10]. In addition, observations optimized for detection of potentially hazardous asteroids will be performed. As a spin-off, these observations will generate a very deep, very broad-band image concentrated around the ecliptic plane.

In Pan-STARRS' standard operation mode, four simultaneous exposures in the same filter will be taken. This will allow very efficient rejection of cosmic rays and other artifacts. In the interests of temporal sampling, and also to provide large numbers of images to ameliorate systematic effects, integration times will be chosen such that the sky noise variance exceeds the read noise variance by factor 10. These integration times range from tens of seconds

[10]See `http://pan-starrs.ifa.hawaii.edu/project/reviews/PreCoDR/documents/PSCoDD_1_4_design_reference_mission.pdf` for the point source detection limits for the various surveys.

in the red to a couple of minutes in the blue. The telescope scheduler will always attempt to obtain observations in pairs separated by a 'transient time interval' on the order of tens of minutes; this allows a clean separation of stationary transients and asteroids.

In order to generate accurate sky-plane surface-brightness images, Pan-STARRS will require dense, all-sky astro- and photo-metric reference stars. These reference catalogs will be generated using the first prototype telescope (PS1).

C.2. Data Pipeline

The Pan-STARRS image processing pipeline (IPP) will perform the following tasks: It will debias and divide out detector response and subtract air-glow and other foregrounds from the raw images delivered from the summit. It will detect objects (primarily point sources) in these processed images to provide instrumental positions and magnitudes. Comparison of these catalogs with external reference catalogs will provide definitive transformation from instrumental quantities (pixel location and ADU counts) to sky coordinates and surface brightness. It will warp the exposures to sky coordinates and combine the four exposures with cosmic ray rejection to make the 'current image'. PSF-matched static sky template subtraction will result in a difference image. Streaks from artificial satellites will be identified and excised. Objects detected from the difference images will be correlated against static sky catalogs and with co-spatial recent difference image detections and will thereby be flagged as stationary variable, stationary transient and moving objects. Finally, after removal of transients, the current image will be accumulated into static sky image.

C.2.1. Moving Object Pipeline

The principal moving objects that Pan-STARRS will be interested in are asteroids. The data from even the first Pan-STARRS telescope will place a heavy load on existing infrastructure (i.e., the Minor Planet Center and related operations) for linking multiple detections of a given moving object, and determining its orbit. Pan-STARRS will maintain a huge and growing database of low-significance ($> 3\sigma$) detections for "pre-covery". It is impractical to export these data, so it is necessary that the linkage and orbit determination processes — currently handled by the Minor Planet Center — be tightly linked into the Pan-STARRS data archive. Thus the moving object pipeline should be capable of importing detections from other surveys, such as LINEAR and DCT. Long term integrations and risk evaluation/announcements should be external to the project.

C.2.2. Basic Data Products

Pan-STARRS basic data products are

- **Instrumental catalogs.** These contain instrumental magnitudes and coordinates for brighter objects detected in the sky-subtracted and calibrated source images. These catalogs are used internally by the IPP, but are also a valuable science resource for precision astrometric and photometric calibration of overlapping external data sets. Postage stamps will also be saved for sufficiently bright objects.
- **Cumulative static sky images.** These consist of signal and exposure maps in sky coordinates (projections of α, δ onto a set of tangent planes that cover the entire sky). These images will be sampled at a pixel scale of $0''.2$. Various versions of these will be maintained, including a 'best' image, to be used as a template, and a 'working' version into which the current images will be accumulated. Various intermediate accumulations

will also be saved, but may be stored on media with higher latency. Images will be convolved with the PSF before accumulation, resulting in optimally weighted combination of good and bad seeing images.

- **Static sky catalogs.** These will be generated periodically from the accumulated static sky image and will include time history of magnitudes measured from the current images. These catalogs will contain extended objects that will have considerably more attributes than the instrumental or difference image catalogs (which are primarily point sources or short streaks).
- **Difference image detection catalogs.** This database will grow continuously with time and will include objects detected at a low significance level (3σ) in order to allow 'pre-covery' of moving objects.
- **Recent source and difference images.** The former will allow one to back-out any erroneously accumulated images in the event of problems. The latter are provided in order to provide, for example, information about precursors to transient events.

Additional data products:

- All sky astro- and photo-metric standards. These will be generated during a 'pre-survey' to be carried out with PS1.
- Moving objects database consisting of detections, putative linkages, and orbits.
- Engineering meta-data.

C.3. Performance Metrics

We now present a detailed discussion of the performance metric for Pan-STARRS. We first present an analysis of the image quality factor $1/\Delta\Omega$ appearing in the faint point source detection figure of merit.

C.3.1. Image Quality Analysis

The generally agreed primary figure of merit for survey telescopes is the étendue $A\Omega$ multiplied by the inverse PSF area $\Delta\Omega$. The former is trivial to estimate (Pan-STARRS has an $A\Omega$ that is approximately 5 times smaller than the proposed 8.4m LSST design). The latter is more subtle and should be computed as the zeroth moment of the square of the optical transfer function (i.e. it is the bandwidth of the imaging system). The figure of merit defined this way converts directly to the integration time required to detect a faint point source against the air-glow.

In general, we expect Pan-STARRS image quality to suffer as a result of several effects. First, there is its relatively large finite pixel size ($\sim 0.3''$) and detector smearing. Second, diffraction effects will be larger for a smaller pupil, and will tend to become important at long wavelengths as the Fried length r_0 becomes comparable to the aperture size. On the other hand, the improvement in the PSF arising from guiding is also greater for a small D/r_0, so these competing effects need to be considered together.

Figure 32 shows the figure of merit, relative to that for a natural seeing image, as a function of filter passband and also over a plausible range of seeing ($0''.4 <$ FWHM $< 1''.6$ at $\lambda = 0.5$ micron). The dashed line shows the FOM for natural seeing, which scales as FWHM^{-2}. The lower of each group of solid lines shows the combined degradation from finite pixel size, charge-diffusion and diffraction effects. These curves were computed assuming a circularly symmetric pupil with a central obscuration 0.6 times the size of the primary mirror. In extremely good

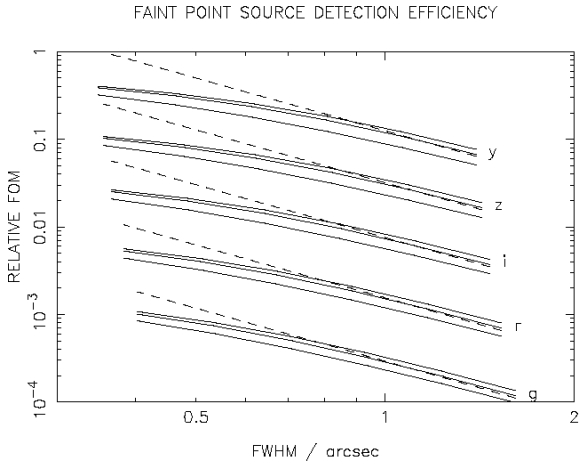

Fig. 32.— Image quality FOM as a function of the natural seeing width. The image quality is obtained by taking the atmospheric OTF, corrected for fast guiding, multiplying by the pupil and detector (pixel and charge diffusion) OTFs and then integrating.

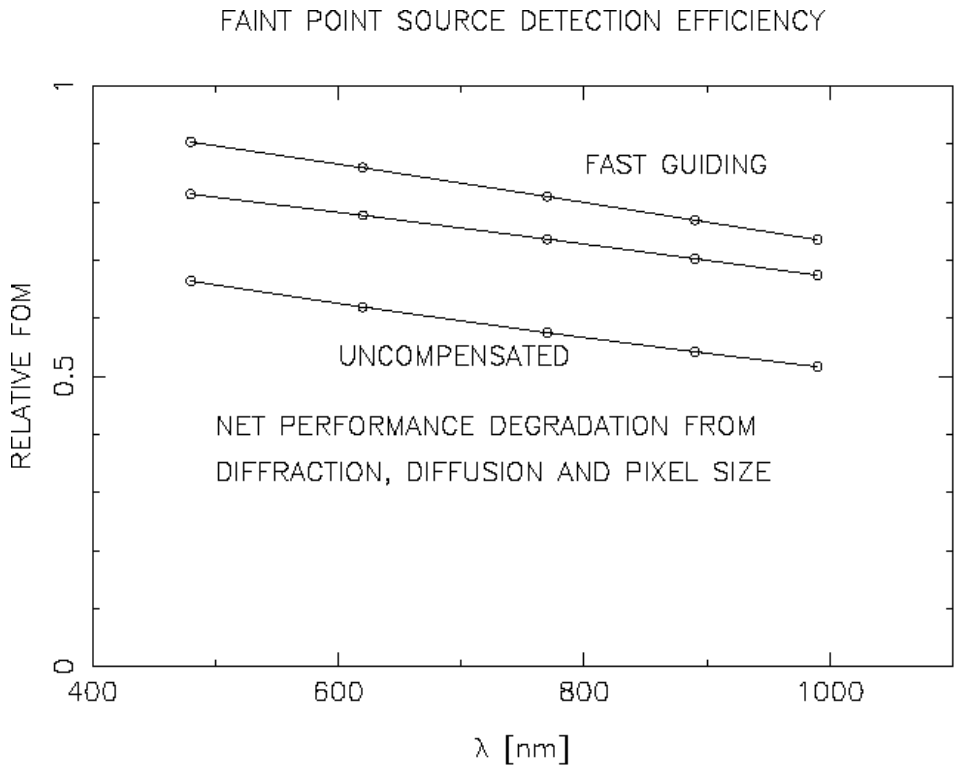

Fig. 33.— Image quality metric as a function of passband, computed by averaging the FOM shown in Figure 32 weighted according to the seeing distribution model shown in Figure 34.

seeing conditions, the degradation is substantial, particularly at long wavelengths. The upper solid curves show the relative FOM assuming fast guiding with an outer scales of 20 meters, 50 meters and infinity respectively. The image improvement from fast guiding suggested here is optimistic in several respects; it assumes that all of the seeing arises at very low altitude (a few hundred metres or less), and if only some fraction of the wavefront deformation variance arises from these levels the gain in performance is reduced proportionately. While we do not have direct measurements of the outer scale, measurements at La Silla and Paranal show outer scales of around 20 meters or less. And, on this basis, our best estimate of the final image quality FOM is that it will lie somewhere between the two lower solid lines in each group.

These calculations assume a single charge diffusion scale. More realistically, we should allow for a decrease at long wavelengths where silicon becomes transparent. However, the gain is quite small as diffraction is then becoming the dominant factor.

Figure 33 shows the net performance degradation, as a function of wavelength, averaged over the distribution of r_0 values corresponding to the distribution of seeing disk widths shown in figure 34. In summary we find that the net performance degradation ranges from about a factor 0.7 in the blue to a factor 0.6 in the near IR.

C.3.2. Read Noise and Read/Slew Times

For a given integration time, a slower telescope will collect fewer photo-electrons and so will have a relatively larger component of read-noise. This introduces a factor $(1 + \sigma^2_{\rm read}/N_{\rm photons})$ in the overall FOM. As discussed, integration times will be chosen such that this penalty is relatively minor. In dark conditions, and at zenith, the broad solar system filter results in about 12 photo-electrons per pixel per second. In a thirty second exposure will give 360 electrons from the sky, so a read noise of $4e^-$ rms would impose a penalty of only a few percent.

Finally, for Pan-STARRS, there is an overhead for read time and slew/settle time. Assuming 3 second read followed by 3 seconds slewing and settling, there is a 80% duty cycle. For a large telescope design like the 8.4m LSST the overhead from slewing and settling may be significantly higher.

C.3.3. Collision Hazard Reduction

We have performed detailed simulations to establish the performance of Pan-STARRS for detecting potentially hazardous objects. A set of PHA orbits were drawn from the empirical model of Bottke *et al.* (Icarus, 2002) and the

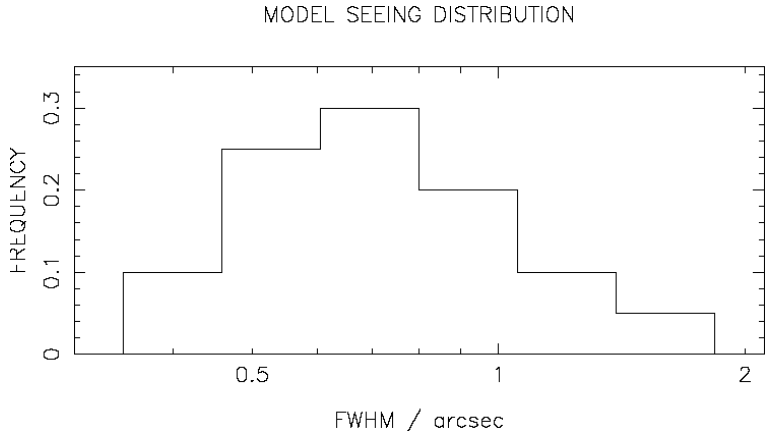

Fig. 34.— Model for the seeing distribution used in computing the image quality FOM in Figure 33.

objects were assigned a collision risk based on their orbital period and collision velocity. The objects were then propagated forward for a period of 20 years, and, for each night, the observability of each object was calculated and, if observable, the net collision hazard was appropriately reduced. The observability calculation takes into account the air-mass (and consequent increase in seeing), sky background (including Zodiacal light), trailing losses, seeing variation, weather fluctuations etc. For details, see `http://pan-starrs.ifa.hawaii.edu/project/people/kaiser/neos/collis`

The resulting reduction in collision hazard is shown in Figure 35 for a variety of collision energies and for current surveys (modeled as having a detection limit $m_{\rm lim} = 21$), Pan-STARRS ($m_{\rm lim} = 24$) and a larger future instrument like the proposed 8.4m LSST with $m_{\rm lim} = 25$, but starting operation six years after Pan-STARRS. This plot shows the current surveys becoming near complete for km size objects (30,000 megaton impact energy; MT) by the time Pan-STARRS becomes operational. For $\sim$ 1000 MT impactors ($H' \simeq 21$), current surveys will remain only partially complete, and will be rapidly overtaken by Pan-STARRS. Over the 10 year project lifetime, Pan-STARRS will be able to eliminate more than 80% of the collision risk for these energies. The bottom panel shows the completeness for smaller objects (140 MT) where Pan-STARRS remains less than 50% complete even after 20 years of operation. For such energies, Pan-STARRS would eventually be overtaken by an 8.4m LSST, but the completeness is still only partial even after 20 yrs of operation.

In summary, Pan-STARRS, with its large increase in sensitivity over current instruments, will be most effective for detecting PHAs down to $\sim$ 1000 MT impact energies. For smaller objects ($\sim$ 100 MT) neither Pan-STARRS nor a 8.4m LSST class instrument will approach full completeness over any reasonable operation time.

C.4. Pan-STARRS Project History

A proposal was submitted by the Institute for Astronomy (IfA) of the University of Hawaii (UH) in early 2002 in response to an Air Force Broad Area Announcement inviting bids to develop observatory technology.

First year funding arrived in September 2002, and the first year activities were concerned primarily with developing the design and ramping up the man-power in management and engineering required to carry off this project. That ramp up is now (i.e., in April 2004) nearly complete. After preliminary internal requirements review in August 2003, the project is now embarking on a rapid process of requirements and design reviews for the various components. The optical design has passed preliminary design review, and contracts for design studies by telescope and optics vendors have been initiated. Design of a range of test devices of various types has been completed, and a fabrication run of 24 15 cm wafers has been initiated and processing is 50% complete.

Sites on both Mauna Kea and Haleakala are being considered, and site testing is being performed in parallel at both locations.

In December 2003, it was decided by the IfA, the Air Force, and the project to develop a first prototype telescope 'PS1' to be deployed in what is currently the LURE lunar ranging experiment building on Haleakala. PS1 will be a full-scale telescope with a full focal plane and is intended to stimulate the development of and to 'shake-down' the numerous software and hardware sub-systems and to allow integration of these sub-systems prior to deployment of the full Pan-STARRS array.

First light for PS1 is scheduled for January 2006, with deployment of the full array within a further two years.

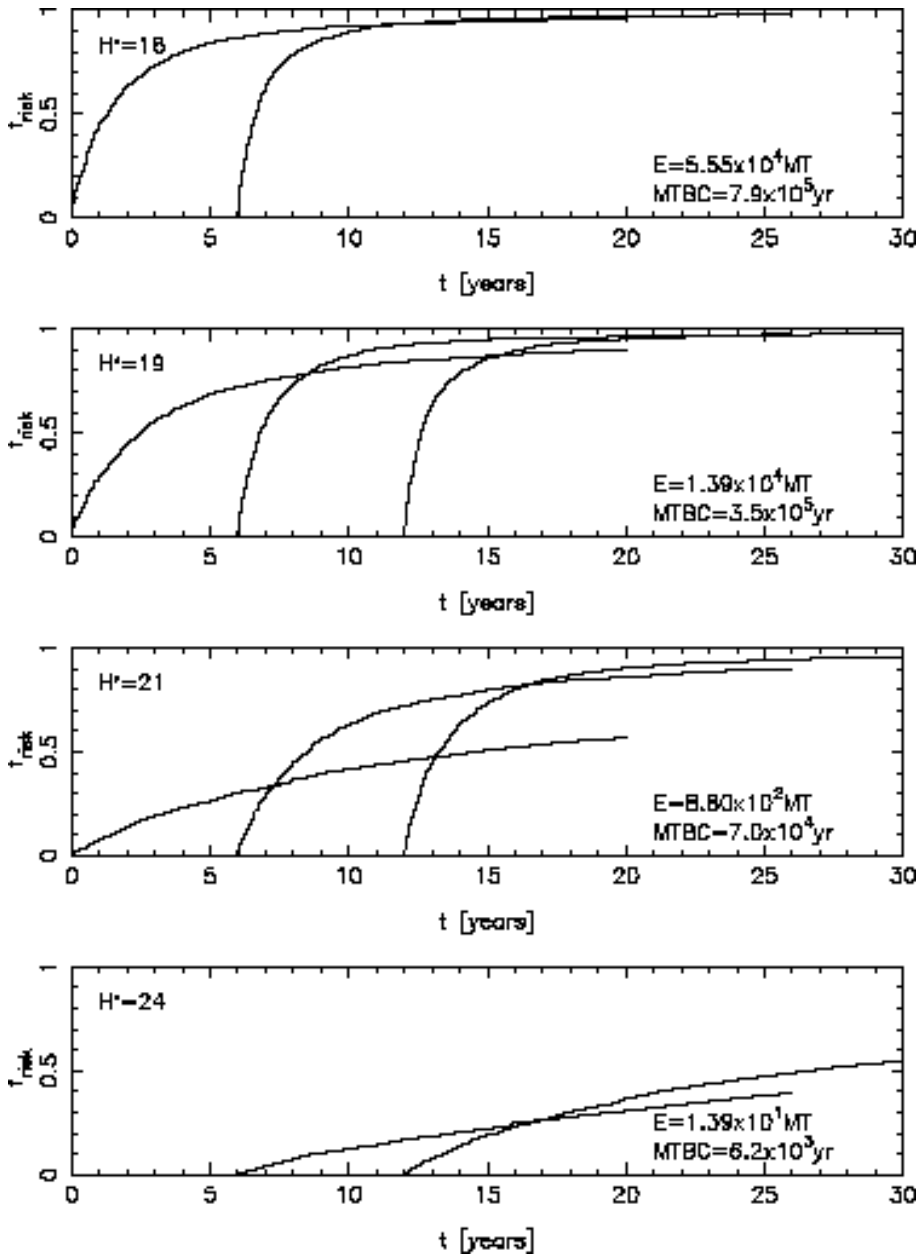

Fig. 35.— Completeness of reduction of collision hazard for various collision energies (by panel) and for detection limits of $m_{\rm lim} = 21, 24, 25$ (curves from left to right). Note that the start time of the different assumed detection limits are different.

D. The Membership and Affiliation of the LSST Science Working Group

The LSST Science Working Group was convened by Jeremy Mould, director of NOAO, under the aegis of a request from the National Science Foundation. The membership of the group is as follows, together with their principal areas of scientific and technical expertise.

Gary Bernstein, University of Pennsylvania, garyb@physics.upenn.edu (weak lensing, KBOs, wide-field imaging)

Andy Connolly, University of Pittsburgh, ajc@tiamat.phyast.pitt.edu (Photometric redshifts, galaxy clustering, data systems)

Kem Cook, Lawrence Livermore National Laboratory, kcook@igpp.llnl.gov (wide-field imaging; microlensing)

Daniel Eisenstein, University of Arizona, eisenste@as.arizona.edu (galaxy evolution, cosmology, photometric calibration)

Peter Garnavich, University of Notre Dame, pgarnavi@miranda.phys.nd.edu (supernovae)

Alan Harris, Jet Propulsion Laboratory, harrisaw@colorado.edu (NEOs and asteroid science)

Fiona Harrison, California Institute of Technology, fiona@srl.caltech.edu (the variable sky)

Željko Ivezić, Princeton University/University of Washington ivezic@astro.princeton.edu (asteroids, surveys, image processing, Galactic structure, the variable universe)

Dave Jewitt, University of Hawaii, jewitt@ifa.hawaii.edu (KBOs and asteroid science)

Nick Kaiser, University of Hawaii, kaiser@ifa.hawaii.edu (weak lensing, cosmology, wide-field imaging, image processing, Pan-STARRS Project Scientist)

Steve Larson, University of Arizona/Lunar and Planetary Lab, slarson@lpl.arizona.edu (asteroid science)

Dave Monet, US Naval Observatory/Flagstaff, dgm@nofs.navy.mil (stellar science; astrometry)

David Morrison, NASA Ames, dmorrison@arc.nasa.gov (NEOs and asteroid science)

Mike Shara, American Museum of Natural History, mshara@amnh.org (stellar science, the variable universe, public outreach)

Alan Stern, Southwest Research Institute, astern@boulder.swri.edu (KBOs, solar system science)

Michael Strauss (Chair), Princeton University, strauss@astro.princeton.edu (quasars, large-scale structure, surveys)

Chris Stubbs, Harvard University, stubbs@physics.harvard.edu (weak lensing, wide-field imagers, supernovae and variable objects, data processing)

Tony Tyson, University of California, Davis, tyson@lucent.com (weak lensing, wide-field imagers, image processing, 8.4m LSST director)

Dennis Zaritsky, U. Arizona/Steward Observatory, dzaritsky@as.arizona.edu (galaxy properties and photometry, wide-field imaging)

While the SWG included no membership from the scientific staff of NOAO, we benefitted greatly from input from a variety of people, including, but not restricted to: Sidney Wolff, Jeremy Mould, Abi Saha, Knut Olsen, Chuck Claver, Richard Green, Beatrice Muller, and Chris Smith, as well as Philip Pinto of the Department of Physics at the University of Arizona.