# Pan-STARRS — A Large Synoptic Survey Telescope Array


Nick Kaiser[a], Herve Aussel[a], Barry Burke[b], Hans Boesgaard[a], Ken Chambers[a], Mark Chun[c], Jim Heasley[a], Klaus Hodapp[c], Bobby Hunt[d], Robert Jedicke[f] Dave Jewitt[a], Rolf Kudritzki[a], Gerry Luppino[a], Mike Maberry[e] Eugene Magnier[g], Dave Monet[h], Peter Onaka[a], Andrew Pickles[c], Pui Hin Rhoads[a], Theodore Simon[a], Alex Szalay[i], Istvan Szapudi[a], Dave Tholen[a], John Tonry[a], Mark Waterson[e] and John Wick[j]

[a]Institute for Astronomy, 2680 Woodlawn Road, Honolulu, Hawaii 96822, USA
[b]MIT Lincoln Laboratories
[c]IfA Hilo Facility
[d]Maui High Performance Computing Center
[e]IfA Waiakoa Laboratory, Maui
[f]Lunar and Planetary Laboratory, U. Arizona
[g]Canada-France-Hawaii Telescope Corp.
[h]USNO Flagstaff Station
[i]Physics and Astronomy, Johns Hopkins University
[j]Science Applications International Corp, Tucson



## ABSTRACT

The IfA and collaborators are embarking on a project to develop a 4-telescope synoptic survey instrument. While somewhat smaller than the 6.5m class telescope envisaged by the decadal review in their proposal for a LSST, this facility will nonetheless be able to accomplish many of the LSST science goals. In this paper we will describe the motivation for a 'distributed aperture' approach for the LSST, the current concept for Pan-STARRS — a pilot project for the LSST proper — and its performance goals and science reach. We will also discuss how the facility may be expanded.



**Keywords:** Wide-field imaging, Survey telescopes, Large-format CCDs, Near earth asteroids.


## 1. INTRODUCTION

Wide field imaging surveys are undergoing a renaissance, led by advances in detector technology, coupled with enabling technology in readout electronics and data processing capability. Examples of recent projects which exploit these developments are the Sloan Digital Sky Survey (SDSS) in the optical and the 2 micron All Sky Survey (2MASS) in the near infra-red.

The National Academy of Science's Decadal Review of Astronomy (1) recognized that the time is ripe for for a further leap in wide-field optical imaging and proposed as a major priority the development of a Large Synoptic Survey Telescope (LSST). The decadal reviewers envisaged a 6m class telescope with a very wide field of view ($\simeq 3^\circ$ diameter or 7 square degrees). The requirements set out were that such an instrument should be able to survey the entire visible sky to a detection limit $m_R \simeq 24$ in a week or less. Such an instrument would generate two basic data products. By scanning the sky multiple times over its operational life it would be able to generate an accumulated image of the static sky of unprecedented depth. In addition, new images would be PSF matched with, and subtracted from, the static sky image to produce a stream of difference images to reveal transient, variable or moving objects. With this combination of deep cumulative images and unique time resolution such a facility would be able to to fulfill a multitude of science goals (see table 1).


For further information visit http://poi.ifa.hawaii.edu or contact kaiser@hawaii.edu.

**Table 1.** Science goals of a large synoptic survey instrument.

| | |
|---|---|
| Solar system | NEOs — KBOs — comets |
| Extra-solar planet searches | transits — micro-lensing |
| The galaxy | parallax survey — proper motions — galactic structure |
| Supernovae | expansion history of the universe — cosmological parameters |
| Galaxy clustering | shells at $z \sim 1$ from photo-$Z$s — feeder for spectroscopy |
| Weak lensing | dark matter — power spectrum — cosmology |
| Variable stars | distances |
| Transients | GRBs — lensed quasars — new phenomena |
| Active galactic nuclei | evolution — host galaxy environment |

The LSST proposal did not emerge from a vacuum. Prior to the Decadal Review, Tony Tyson and Roger Angel had campaigned vigorously for community support for a "Dark Matter Telescope" (2; 3), which has its roots in Tony's long standing dreams of a dedicated telescope for weak lensing. In the wake of congressional mandates for NASA to attack the problem of Near Earth Asteroids (NEOs), this goal was taken on board, as it was realized that such a telescope could profitably attack both objectives. The DMT proposal is for an 8.4m telescope (effective collecting area roughly equivalent to a 6.5m filled aperture), and is a very fast 3-mirror design similar to Willstrop's 'Mersenne-Schmidt' design (4).

The DMT is an exciting concept, and would deliver an *etendue* — the product of collecting area and field of view — much greater than any existing facility and also larger than some planned projects such as the UK VISTA telescope (5). However, there is an alternative, which is to use a 'distributed aperture' approach; i.e. to use an array of smaller telescopes, and to combine the resulting images incoherently in the computer. This may seem rather odd, in the face of moves to go beyond the current generation of $8 - 10$m class telescopes to 30m or even larger. The driving goal for such instruments, however, is to couple them with full adaptive optics (AO) wave-front correction in order to deliver diffraction limited images. Unfortunately, full AO is impossible for very wide field imaging because of the 'isoplanatic angle' problem, and consequently any LSST will have to accept the natural seeing. Without full AO, the motivation for a large single aperture is removed and it becomes a question of economics, engineering and politics whether it is better to use a single telescope or an array of smaller instruments. Whether to use a single telescope or an array for the LSST is a question that the community will need to address. To inform this discussion, in section §2 we will review the advantages and disadvantages of the two approaches. The main advantages of distributed apertures for many applications are low cost, speed of construction, and low risk. On the negative side, there is the additional cost associated with the necessary multiplexing of detectors, read-out electronics etc., though, as we shall see, the economics here are not quite as simple as one might think. In §3 we will describe the current benchmark design for the Panoramic Survey Telescope and Rapid Response System (Pan-STARRS) project; this is a somewhat smaller scale instrument than envisaged by the decadal reviewers for the LSST. Pan-STARRS can be thought of as a pilot project for a full-scale LSST project using an larger array of telescopes. Nonetheless, as we shall show, the instrument is capable of delivering much of the science promised by a LSST. In §4 we will describe the power of Pan-STARRS for detection of potentially dangerous NEOs. Finally, in §5, we will discuss the future expansion of the Pan-STARRS project.

## 2. SINGLE *VS.* DISTRIBUTED APERTURES.

We now discuss the relative merits of single *vs.* distributed apertures for wide-field surveys. For many applications the appropriate figure of merit for such instruments is the *etendue*, though, as we will note, there are exceptions. The discussion below will not assume any specific number of telescopes in the distributed aperture array. Pan-STARRS is currently planned to be a small array of 4 telescopes of $D \simeq 1.8$m, but the design is scalable, and some of the points raised below are relevant for larger future systems which would be directly comparable to, and of similar expense as, the DMT. We will start with features where the distributed aperture approach seems to offer substantial advantages and we will conclude with a discussion of the cost penalty for duplication of detectors.

**Telescopes:** It is generally agreed that the cost of telescopes rises faster than the collecting area. If the figure of merit is simply the *etendue* then small telescopes offer a financial cost advantage. Small telescopes ($D \lesssim 2$m say) can also be constructed quite rapidly, and this offers an opportunity cost advantage. This is particularly important in view the analog of Moore's law for detectors which is driving the costs of detectors down rapidly. Small telescopes can exploit optically slow and relatively low-risk designs, while large designs like the DMT are driven to be very fast. Fast designs have problems with e.g. interference filters and stiffness, the latter being a high priority for the LSST since to cover the whole sky rapidly requires relatively short exposure and slewing/settling times. Arrays of small telescopes can be housed in low profile enclosures with relatively low environmental impact; this is a serious consideration for sites such as Mauna Kea and, no doubt, elsewhere. Finally, fast designs require more stringent tolerance on alignment of the optical components and on e.g. flatness of the focal plane.

**Number of images:** Typically, an array of small telescopes will collect a larger number of images of any field than a single large telescope. If the individual integration times are held fixed then an array of $N_t$ telescopes (with the same total collecting area as the single large telescope) will need to take $N_t$ times as many images to reach a given depth for the static sky images, and the same is true for measurement of any slowly varying phenomena (time-scale much greater than the interval over which the images are taken). For relatively rapid transient phenomena, however, there are some advantages to having large numbers of shallow exposures than small numbers of deep images. One advantage is removal of backgrounds from cosmic rays etc. If we have say a 5-sigma detection in a single image from a large telescope then it is hard to know if it is real, whereas if the image is formed from $N_t$ images from an array any artefact will appear as a $5\sqrt{N_t}$-sigma detection in just one of the contributing images. Interference fringe patterns also tend to average out better with more exposures, and similarly for variable sky transparency. Combining multiple images also decreases the effective pixel scale. This allows, for instance, anti-aliasing to be applied, which is helpful if deconvolution is to be attempted. An array of telescopes can be used in a mode in which the same field is imaged simultaneously in multiple passbands. How much of an advantage this gives depends, crucially, on the time-scale for the phenomena that one is trying to measure. One might argue that a large-telescope could take much shorter exposures; however, there is a cost, since if the read-out overhead is not to be increased this will require more investment in read-out electronics.

**Dynamic Range:** Averaging many images from small telescopes allows one to measure properties of bright stars faithfully. This is useful for astrometric and photometric calibration. The low dynamic range problem for large telescopes may, however, be ameliorated with the advent of the kind of detectors we describe below.

**Temporal sampling:** Distributing the apertures in longitude — say at Mauna Kea and the Canary Islands — would allow better sampling of light-curves of variable sources.

**Averaging of atmospheric field distortion:** In setting up a system of astrometric standards for a large synoptic survey the limiting precision comes not from photon counting and telescope resolution, but from the atmosphere. This produces random distortion of the field, the large-scale components of which have a long time-scale and so average out rather poorly. In a given time set aside for developing an astrometric standard system an array of small telescopes will be able to average over many more realizations of this distortion, and so will have better absolute astrometric precision.

**Wide FOV operations:** A major advantage of distributed apertures is the ability, if desired, to operate in wide field of view mode, with either all or groups of the telescopes pointed at independent fields. This is advantageous for detection of short time-scale transient events; if the cumulative distribution of counts with flux density $N(>F)$ has logarithmic slope $-d\ln N/d\ln F < 2$, as is the case for most known populations, then a wide shallow survey with an array of small telescopes will detect objects more rapidly than a large single telescope with the same aperture. Another application is detection of smaller NEOs on their way to impact the Earth (see below). These are both applications where the *etendue* is not the relevant figure of merit.

**Low order adaptive optics correction:** We have said that AO is impossible over wide fields. An exception is the WFHRI concept (6), which can be implemented with an array of small telescopes. This would give images with $\sim 30\%$ of the light in a diffraction limited core, and which would give considerable gain, particularly for astrometry and shape-measurement of small sources. The WFHRI concept, however, requires a much finer pixel scale than envisaged for e.g. the Pan-STARRS project.

**Cost of detectors:** An obvious disadvantage of the distributed aperture approach is the required multiplexing of detectors. However, one should not imagine that this simply multiplies the detector cost component by the number of telescopes in the array. This is because, in large designs, the pixel size is chosen not so much from detector considerations as from the requirement to match the plate scale delivered by large telescopes. The DMT Willstrop-style design, for example, can be scaled down to about 4m primary aperture without changing the $f$-ratio, and consequently without any increase in the cost of the imaging array (the cost of the CCDs scaling primarily as the area of silicon rather than as the number of pixels). At this point the smearing of images by charge-diffusion and finite pixel size starts to become important. To go to smaller apertures without penalty requires adoption of slower designs to keep the PSF of the telescope matched to the resolution deliverable by the CCD. Scaling from the cost of recent detector systems, one might still conclude that the extra detector costs would outweigh the advantages of a distributed aperture for $D \lesssim 2$m. As we will argue below, it is possible to obtain a large decrease in the price of detectors. These savings come mainly from 3 factors: *1)* It is possible to increase the yield considerably by making detectors in which each monolithic unit is in fact an array of smaller independent devices, or 'cells', and to live with the small number of cells that will be rendered unusable due to manufacturing failures. *2)* The use of consumer-off-the-shelf (COTS) components developed for HDTV in the read-out electronics; development of such systems is now proceeding in parallel at several laboratories. *3)* Reduction of design and development costs (per camera).

**Read-noise:** Another disadvantage for very small telescopes is that it takes longer integration time to become sky-noise, rather than read-noise, dominated. For telescopes of $\sim 1.5 - 2$m diameter this is not a problem for integration times $\gtrsim 10$s.

## 3. THE PAN-STARRS PROJECT

The Institute for Astronomy at the University of Hawaii, together with collaborators SAIC and MHPCC, have obtained support for the development of the Pan-STARRS system. This project has two major goals: The first is to deliver much of the science and utility promised by the LSST, but on a faster time-scale. The second is to serve as a pilot-project to test the feasibility of the distributed aperture approach to the LSST and to demonstrate the ability to deliver detector systems at the required cost.

The planned time-scale for deployment of the Pan-STARRS system is 4 years. The first year will be devoted to development of the full design specification for the system, followed by three years of construction, software development etc., leading to integration of the system and first light in 2006. We will now describe the 'benchmark' design for the various parts of the system, though it is anticipated that the various design parameters may evolve during the design phase.

### 3.1. Bench-mark Design for Pan-STARRS

The Pan-STARRS project has three major components: the optical system, the detectors and read-out electronics, and the data-processing pipeline and data-archiving and distribution system.

#### 3.1.1. Optical system

The bench-mark design calls for 4 telescopes, each of 1.8m primary diameter and with fairly conventional Ritchey-Chretien design with a wide field corrector and delivering a 3-degree diameter (7 square degree) field of view. Similar instruments have been used widely in wide-field imaging — a recent example being the somewhat larger 2.5m SDSS telescope — and there are various suppliers capable of delivering such systems on a time-scale $\sim 2$ years. The desired plate-scale is $\simeq 30\mu$m/arcsec which requires an $f/3.3$ system.

Several Hawaii sites are being considered. The leading candidate is to house the telescopes in an extension of the UH 2.2m telescope building. This currently houses a defunct Coude spectrograph. Also under consideration are the UH 24" facility and sites on Haleakala. The final site decision will be made based on the result of site testing, and consideration will also be given to data communication requirements.

**Figure 1.** Artist's impression of the Pan-STARRS system.

### 3.1.2. Detector system

The detector system is planned to have an angular pixel size of $0''.3$, which is well matched to the seeing at Mauna Kea. At the plate-scale above, this requires $\simeq 1.0 \times 10^9$ pixels per camera and a pixel size of $10\mu$m. The physical size of the detector is roughly 32cm on a side. The focal plane array is to be constructed as an array of arrays of independently addressable cells. The current benchmark is to have an $8 \times 8$ array of monolithic units of size $\simeq$ 5cm, each of which would consist of an $8 \times 8$ array of $512 \times 512$ pixel cells. As mentioned, a great advantage of this system is that manufacturing defects, which, in more conventional designs, would take out an entire monolithic device, will disable only one cell. We envisage that each 4K $\times$ 4K device will have a few dead cells. By dithering exposures, and making simultaneous imaging, the impact of these dead cells on the science output is negligible, and the resulting increase in yield will reduce the cost of the CCDs considerably.

Our hope is to employ orthogonal transfer (OTCCD) technology (7). In this mode those cells on the detector containing bright stars will be read-out at a rapid rate, and the positions of the stars analyzed to provide guiding information for the cells being used for science integrations. The OTCCD promises an increase in image quality by taking out both atmospheric image motions and telescope wobble. A factor $\sqrt{2}$ decrease in FWHM is feasible. This converts to a factor two increase in system efficiency for point source detection and a factor 4 increase in efficiency for astrometry. See John Tonry's talk in these proceedings for more details.

The read-out electronics chain will employ off-the-shelf components developed for HDTV. Similar systems are being developed at NOAO and at Keck. Such systems will provide great increases in bandwidth over existing systems, and the goal is to be able to read out the entire array in $\sim$ 2s with a read-noise of a few electrons.

### 3.1.3. Data-processing pipeline

In production mode, and assuming simultaneous imaging, the data processing pipeline will perform the following sequence of steps. First the new image data will be read out to a set of buffers. The data for each 4K$\times$4K devices will then be flattened and corrected for instrumental gain variations. Stars and suitably centrally concentrated galaxies will then be detected, and the resulting catalogs will be compared to a reference catalog to obtain a mapping from detector coordinates to sky coordinates. It is well established that such transformations can readily be obtained with $\sim$ 5mas relative astrometric precision. The precision for absolute astrometry will be limited by the precision of the astrometric reference catalog (see below). The cross-linking with the reference catalog will also provide transformations from the instrumental flux density to properly calibrated sky brightness. At this point, the 4 image planes will be mapped to sky coordinates — it is anticipated that the sampling in sky coordinates will be performed at a sampling rate $\sim$ 1.5 times finer than the raw detector pixel size — and combined to generate a contiguous 7deg$^2$ image for each field. By combining 4 images it is possible to almost completely remove backgrounds such as cosmic ray hits and other artifacts. The PSF for such an image — which is expected to be a smoothly varying function of position within the field — will be measured and modeled, and the image will be convolved with its PSF. The accumulated static sky image for the field will then be digitally filtered by Fourier techniques to match the PSF of the new image, and will be subtracted therefrom to provide a difference image.

The resulting stream of difference images, which will subsequently be analyzed for detection of moving, transient or variable objects, is one of the two basic products of the Pan-STARRS pipeline. Finally, pixels in the new image which correspond to transient or moving objects will be flagged as unreliable and the rest will be accumulated, with appropriate weight, into the static sky image. This is the other basic image data product.

The weighted sum of these PSF convolved images can be shown, for low surface brightness sources at least, to provide an accumulated image which is optimal for all purposes. For bright sources — those for which the contribution to the pixel electron count is greater than that from the sky — this way of combining images is not optimal for all applications. It is anticipated that the pixels for such bright source regions — mostly stars — will be stored rather than simply accumulated.

This co-addition algorithm is, of course, only optimal for the static sky. As a goal is to detect variability of objects on many different time-scales, there will in fact be a series of accumulated images produced which give the average over various time-scales ranging from months to years.

It was originally imagined that the read-out time would be on the order of 30s, and that the integration time would be $\sim$ 60s. This would provide an averaged data-rate of about 45Mpix/s, or about 350Mbit/s (assuming that, with compression, one requires roughly one byte per pixel). With the $\sim$ 2s read-out times now envisaged, it is likely that images will be obtained perhaps once every 30s, which would triple the data communication and processing requirements. It is estimated that, using the kind of processors now available at the $\sim$ \$1K price-range, something on the order of 250 nodes would be required.

We have assumed above that the astro/photometric reference catalog has been generated. We plan to spend a large part of the first year of operation making observations in a mode optimized for the generation of such a catalog. This would involve taking images with quite different spatial sampling than for the survey proper (with fractional field shifts and with rotations of the focal plane) and the results of these observations being processed by separate pipelines for astrometry and photometry.

#### 3.1.4. Data access and archiving

The Pan-STARRS array will collect data at a rate of $\sim 3 - 10$ terabyte (TB) per night. Over several years of operation this amounts to $\sim 3 - 10$ petabytes of data. As already discussed, it is not necessary to keep all of this data on-line; all that are needed to be maintained are the cumulative images and catalogs. The difference image stream, to be useful, needs to be processed in real time for applications such as NEO searches, though compressed versions can be saved at little cost. The all-sky images and catalogs will be made publicly available *via* the world-wide-web. The image data will comprise on the order of 100TB, and the object catalog data perhaps 10TB. The former are rather simply indexed by position on the sky, so it is a relatively minor problem to devise a user interface such that users can request to download patches of this image database. It is impractical to allow distribution of the entire image data *via* the web, so applications that need to access the entire image data (galaxy clustering, weak lensing) will need to be incorporated with the Pan-STARRS system.

The catalog data base, while relatively small in terms of bytes, represents a major challenge in database engineering. Here the indexing of the data is non-trivial, and the optimal indexing depends on what type of questions the science users will ask, and which, of course, is very hard to predict. We plan to support NVO-style queries *via* a standardized interface. Such work is already under way to support the distribution of data from the SDSS, 2MASS and other surveys and the Pan-STARRS project will further develop these interfaces.

## 3.2. Pan-STARRS Performance Goals

### 3.2.1. Photometric performance

Assuming simultaneous imaging and 5-sigma detections, the integration times required for Pan-STARRS to detect a $R = 24$ point source (the goal set out by the decadal reviewers) is

$$t(R = 24.0) = 58\mathrm{s}(\mathrm{FWHM}/0''.6)^2 \tag{1}$$

for an object with $V - R \simeq 0.4$ as is typical for asteroids, the corresponding integration time in the $V$-band is

$$t(V = 24.4) = 67\mathrm{s}(\mathrm{FWHM}/0''.6)^2 \tag{2}$$

and if we use a broad filter, encompassing both the $V$ and $R$ band-passes, the required integration time is

$$t(R + V) = 31\mathrm{s}(\mathrm{FWHM}/0''.6)^2. \tag{3}$$

To obtain these numbers we have used measurements of the sky surface brightness obtained at the CFHT and we have assumed that the overall QE of the system will be similar to CFHT (which has similar optical elements). We have assumed a fiducial image quality appropriate for Mauna Kea.

For wide-field mode — where the telescopes are pointed at different fields — the detection limit for the same integration time is about $2\times$, or $\simeq 0.8$ magnitudes, brighter.

With the pixel size and $f$-ratio assumed above, the night sky will generate $\sim 3\mathrm{e}^-$/pixel/s for the standard broad band filters, or $\sim 6\mathrm{e}^-$/pixel/s for the $R + V$ filter. With a few electrons read-noise this means that for any integration time $t \gtrsim 10$s the images will be strongly sky-noise limited.

The figures above give the absolute statistical limiting sensitivity for these integration times (a five sigma detection as defined here gives a fractional precision for flux density of 20%). Real detection limits may be slightly brighter due to flat-fielding errors etc., but experience shows that it is quite feasible to get to within say $\sim 20\%$ of the theoretical limits. For bright objects (pixel values greater than the sky background) the fractional precision for flux density measurements is just $1/\sqrt{N_\gamma}$ with $N_\gamma$ the number of source photons detected. The absolute photometric precision for bright objects may, however, be compromised by such factors as patchy sky-transparency fluctuations and also intra-pixel sensitivity variation.

### 3.2.2. Astrometric precision

The limiting precision for location of faint point sources (1-sigma error per component) is

$$\sigma_x \simeq 70\mathrm{mas}\left(\frac{5}{\nu}\right)\left(\frac{\mathrm{FWHM}}{0''.6}\right) \tag{4}$$

where $\nu$ is the signal/noise ratio for the detection, and we have used a fiducial 5-sigma detection. For bright sources the statistical limit is $\sigma_x \simeq 250\mathrm{mas}(\mathrm{FWHM}/0''.6)/\sqrt{N_\gamma}$, which can be exquisitely precise. However, there are systematic effects that will likely compromise this.

One such 'systematic' arises from the atmosphere. As is well known, turbulent mixing of air with inhomogeneous entropy or water vapor content in the atmosphere results in tilting of the wave-fronts. The corrugation of the wave front on scales of order 20cm is what accounts for the broadening of the images of point sources. There are, however, also components of the wave-front tilt which are coherent over much larger scales. These have a relatively long time-scale, and consequently do not average out very well in short exposures. If we adopt the conventional view that the seeing arises in one or more layers above the telescope — for which there is quite strong evidence — and assume that the refractive index fluctuations in these layers are as predicted by Kolmogorov theory on scales less than some outer scale $L_{\rm o}$, then the prediction is that the atmospheric will introduce random residual field distortion with characteristic angular scale

$$\theta \simeq 12' \frac{L_{\rm o}}{20{\rm m}} \frac{5{\rm km}}{H} \tag{5}$$

where $H$ is the height of the refracting layer, and with time-scale

$$\tau \sim 2{\rm s} \frac{L_{\rm o}}{20{\rm m}} \frac{10{\rm m/s}}{v}. \tag{6}$$

The rms amplitude of the deflections, in this model, is

$$\sigma_x \simeq 160{\rm mas} \left(\frac{\rm FWHM}{0''.6}\right)^{5/6} \left(\frac{L_{\rm o}}{20{\rm m}}\right)^{1/3} \left(\frac{v}{10{\rm m/s}}\right)^{-1/2} \left(\frac{\lambda}{5\times 10^{-7}{\rm m}}\right)^{1/6} \left(\frac{t_{\rm exp}}{1{\rm s}}\right)^{-1/2} \tag{7}$$

Interestingly, this is independent of the size of the telescope. Now these deflections, for short exposures ($\sim$ 30s) of the kind envisaged in Pan-STARRS survey mode, are quite large. However, since the deflections are predicted to be coherent over scales of order $10'$, these will be taken out when the images are warped from detector to sky coordinates. The limit on absolute astrometric precision is then determined by how long we can afford to integrate on each field during the pre-survey phase where we are setting up the astrometric standard catalog. It is reasonable to assume that we can afford to spend $\sim$ 600s on each field, in which case the simple model above would predict residual errors on the order of 10mas. However, the inputs to the above formula such as the outer scale are not very well measured, and almost certainly fluctuate quite strongly, so it may well be that in many parts of the sky the residuals might be quite large. As mentioned, it is an advantage of the distributed aperture approach that, for the purpose of setting up the astrometric catalog, we could work in wide field mode, in which case the affordable integration time per field is increased, and the absolute astrometric precision correspondingly reduced.

Another limit on absolute astrometric precision comes from how well we can model flexure of the detector etc., but this is harder to predict.

## 4. NEO DETECTION WITH PAN-STARRS

A major selling point for the LSST, and therefore also for Pan-STARRS, is the ability to detect 'killer asteroids' or 'near earth objects' (NEOs). It is this driver that dictates a survey which can cover the available sky to deep detection limits several times per lunation. This, and the fact that NEOs are moving quite rapidly, also dictate the relatively short exposure times. Most of the other science drivers could live with much longer integration times and period between revisits. To a first approximation then, the other science goals can ride on the coat tails of the NEO search survey (there are other factors, such as choice of pass-bands, that impact the various goals, so in reality there will be compromises to be made in order to maximize the science and utility of the project).

Pan-STARRS will be an exceptionally powerful instrument for the detection of NEOs. As mentioned, with the high image quality provided by sites such as Mauna-Kea, this instrument can detect a point source of 24th magnitude in about 30s integration time (60s for standard passbands). The current surveys of NEOs are approximately 50% complete down to a size of about 1km diameter. Pan-STARRS will be able to push this size limit down to about 300m diameter. Simulations show that such asteroids, when they appear at our nominal detection limit of $R = 24$ have angular speeds of about $0^\circ.4$ degrees per day, or about $1''.0$ per minute. In a 30s exposure then such objects will be trailed by less than one FWHM. This is very different from all existing surveys where the sensitivity is lower, and objects of this size are only detectable when they get quite close, and are moving rapidly. For such surveys 'trailing losses' are very important, and have a huge impact on the detection efficiency. For Pan-STARRS, and with $30-60$s exposures, these losses are quite minor. One can estimate the additional integration time required to detect an object moving at $0^\circ.4$deg/day as compared to a stationary point source. It is a $\simeq 20\%$ penalty if the source moves say 1 FWHM of the PSF during an exposure. However, having slightly trailed images can help in making the linkage between objects. It has been argued by proponents of the DMT that exposure times closer to 10s are required in order to avoid trailing losses for 300m sized objects. This seems to us to be overly pessimistic. To be sure, objects of this size will sometimes be found moving considerably faster than the speed quoted above. However, this is because they are then much closer than the limiting distance at which they could be detected and will therefore be correspondingly bright and therefore easy to detect even allowing for the considerable trailing losses they would have in $30-60$s exposures.

Devising an optimal survey strategy for detection of NEOs — let alone a strategy which optimally combines NEO detection and the other science goals — is not easy, and will require careful simulations. However, the following considerations suggest that a nearly complete survey to $\simeq 300$m size is quite feasible. Let's assume we make 60s integration and use standard broad band filters. With 2 second read-out we can hope to make $\simeq 500$ such exposures per night, and therefore cover approximately 3000 square degrees each night. If we restrict observations to zenith distance of less than about 45deg then the total sky available from Hawaii is about $30,000\mathrm{deg}^2$. The visible sky on any night, say within 4 hours of opposition for concreteness, is then about $10,000\mathrm{deg}^2$, of which we can observe about 30% in a single night. Thus, in a single dark run, we might first scan the the most northern 30 degrees in declination and then on each subsequent night scan a strip of sky of the same width but shifted down by 10 degrees. Allowing for hits from weather it is likely that we could cover the whole available sky in this manner in one fortnight. Moreover, each patch of sky (aside from the extreme northern and southern strips) will be observed 3 times with typical separation between visits of 24 hours. Since the width of the strip observed each night is much greater than the distance a typical object can move in a few nights, this avoids the 'picket fence' problem (observing field A while an object is in field B and *vice versa*) and with three observations we will obtain an approximate orbit.

In order to obtain these orbits, it is not sufficient simply to detect the objects multiple times; one must also link them up. In existing surveys, the traditional approach has been to make repeated visits on an interval much less than 24hrs in order to facilitate this linkage. We would argue that this is unnecessary. The linkage problem is most pronounced in the ecliptic plane, where any NEOs will be seen projected onto the more numerous main belt objects. Now the nice thing about the main belters, however, is that they live in a very restricted region of phase space. Consequently, their range of angular velocities and angular accelerations is relatively small, and so, given a set of detections on two subsequent nights, the fraction of sky covered by the error circles on night three obtained by projecting forward all plausible pairs from the first two nights is still very small (we are assuming here a peak density on the sky of about 200 objects per square degree at $R \simeq 24$) so the 24hr revisit period seems to be quite reasonable. The inputs needed for this calculation are somewhat uncertain, so there is some slop in this conclusion. However, this is for the most densely populated parts of the sky. It may be that a somewhat shorter revisit period is required very close to the ecliptic plane, but for most of the sky, the 24hr 'cadence' seems to be quite comfortable.

This was assuming 60s integrations. One could make slightly shorter integrations, and this would allow a single scan of the southern region at the start of the dark run and a single scan of the northern region at the end. This would give an additional detection of most objects with a baseline of approximately 10 days.

Given a crude orbit from say three consecutive observations of a faint NEO then using the astrometric precision figures above it is easy to show that it should be relatively easy to link such an object with its appearance in an image taken say a month later (the error ellipsoid grows quadratically with time). This will give quite a precise orbit, allowing linkage with the appearance of the object say a year later and so on. Note that the search strategy requirements will vary with time. When we first start, it will be necessary to work fast and get triple detections of all objects as described above. The next lunation, however, most of these objects are still visible, so a single detection will suffice to update the orbital parameters. The time scale for the objects within our detection limit to be replaced by other objects which were initially more distant, is on the order of years. This suggests that a more profitable 'steady-state' strategy is to concentrate on the 'fresh' strip of sky where new and previously undiscovered asteroids are appearing. Again, detailed simulations are needed to properly compare the relative efficiency for different strategies.

## 5. BEYOND PAN-STARRS

Pan-STARRS will be a powerful tool to attack the various goals listed in table 1. The cost of the mission is on the order of a few tens of millions of dollars, and it will start making useful science observations *ca.* 2006. However, as mentioned, our goals are more ambitious; we see Pan-STARRS as a pilot project to demonstrate the viability of the distributed aperture approach, and to show that this is the best way to perform the LSST proper (the current estimated cost for which is $\simeq$ \$200M). The impression we have obtained from talking to our colleagues is that, if there is any skepticism, it is in our ability to make detectors on the budget and time-scale required — we wholeheartedly concur that this is the greatest area of risk. Whether this skepticism is justified will be known within a couple of years; in year one we will commission and obtain the results from a CCD fabrication run at Lincoln Labs. We also expect to have prototype electonics in place on the same timescale.

Assuming that the results are positive, and given the arguments in favor of distributed apertures presented above, the way forward will be clear. The LSST should be implemented as an array of many tens of small ($D \simeq 1.5 - 2$m) telescopes. This would then open up some interesting possibilities, both in telescope technology and in science reach.

Regarding telescopes, the choice for Pan-STARRS has been deliberately conservative. We appreciate that we have a challenge before us in developing the 'next generation' detectors. However, if we are considering tens of telescopes, there are alternatives that are worth exploring. One possibility is to use very lightweight replicated mirrors using carbon-fibre composite materials. Such telescopes would require active control, but such technology is now quite mature and the costs are not great. An alternative, less ambitious, approach might be to follow the lead of the Germans with their novel "Hexapod Telescope" (see `http://www.astro.ruhr-uni-bochum.de/astro/hpt/index.html`). This telescope has a thin glass meniscus mirror from Zeiss, supported on a lightweight cell, with active mirror control. This technology delivers an order of magnitude decrease in weight per collecting area, and this should provide a substantial decrease in telescope costs.

Another exciting prospect opened up if we consider arrays of tens of telescopes is to extend NEO detection to smaller objects. The conventional approach to NEO detection has been to scan the sky at a relatively slow rate and to detect large objects at great distances. Consider the $\sim$ 1km size objects that are thought to impact the earth perhaps once every 100,000 years. Such collisions are thought to cause "severe regional damage" and have an appreciable impact on the climate (thought they are not sufficiently energetic to cause "nuclear winter"). Pan-STARRS will detect most such objects quite easily — indeed, perhaps half of these objects are already known — but the probability that there exists such an object that will impact the earth in the next 100 years is very small ($P \sim 10^{-3}$). This does not mean that Pan-STARRS in not worthwhile for this life-saving goal. In the unlikely event that we find such an object on a collision course a campaign could be mounted to deflect the object, and save tens or maybe hundreds of millions of lives. Compared to some other governmental interventions to save lives this is quite a bargain. At a few tens of millions of dollars for Pan-STARRS (we don't need to include the cost of the deflection mission unless this costs more that many tens of billions of dollars as the probability we need to mount the campaign is small) the cost per (expected) man year of life saved is on the order of \$100. This is similar to the estimated cost of saving life by enforcing the wearing of seat belts. This in turn is about an order of magnitude more efficaceous than screening for breast cancer and three orders of magnitude more efficient than the enforcement of air bags in cars (8).

While arguably a hightly cost effective way of saving lives in a statistical average sense, it is most likely that Pan-STARRS will not detect a km size object due to hit us in the next 100 years or so — and presumably our more distant descendants can look out for themselves. Much more likely is that the earth will be struck by a smaller object — the 'existence proof' being the $\sim$ 50m sized object that struck Tunguska in 1908. Now to upgrade the LSST to go much smaller is punishingly expensive; even a factor 3 decrease in size means a factor 10 decrease in flux density at a given distance, and, since we are fighting sky background, this requires two orders of magnitude increase in collecting area. However, there is an alternative, which is to detect such objects as they approach us, when they become very bright very rapidly. As an example, an object deflected from 3AU into a plunging radial orbit destined to collide with us would be 1AU distant a mere 40 days before collision (and therefore hard to detect if it is small), but 10 days before impact it would be 0.3AU distant and would have brightened up by about 5 magnitudes. The signature of such objects is clear; they are relatively slowly moving with $\dot{\theta} \simeq 6 \times 10^{-3}(t/\mathrm{day})\mathrm{deg/day}$, and have angular velocity decreasing linearly towards zero. However, we do not get much warning, and they can appear anywhere on the sky. In order to identify such an object we will need to get 3 positions within maybe 5 days. Clearly this requires very rapid scanning of the sky. This is impossible for a single telescope for any reasonable exposure time — one would need to get say 3 detections on each of 3 nights to be sure that the object is real — so even with 20s individual exposures a single telescope would only monitor maybe $3000\mathrm{deg}^2$ at this cadence. Most objects would therefore hit us unawares. With an array of tens of telescopes it is quite feasible to survey say up to $20,000\mathrm{deg}^2$ each night. The detection threshold would be brighter, of course, but this is not a problem since the objects become very bright. Such an instrument would therefore have something like a 50% probability to detect any incoming object of size in excess of say 50m. The detection probability is somewhat more if the array is distributed in longitude and latitude, but if the object comes from the direction of the sun, as did the 100m size object that flew past in June this year, then detection is impossible. This would give maybe 5-10 days warning of such an impact, which at least would allow time for evacuation of coastal regions or whatever is appropriate.

## References


[1] National Research Council, *Astronomy and Astrophysics in the New Millennium*, National Academy Press, Washington, DC, 2001.

[2] R. Angel, M. Lesser, R. Sarlot, and E. Dunham, "Design for an 8-m telescope with a 3 degree field at f/1.25: The dark matter telescope," in *ASP Conf. Ser. 195: Imaging the Universe in Three Dimensions*, pp. 81–022, 2000.

[3] A. Tyson and R. Angel, "The large-aperture synoptic survey telescope," in *New Era of Wide-Field Astronomy*, R. Clowes, A. Adamson, and G. Bromage, eds., pp. ??–??, 2000.

[4] R. V. Willstrop, "The mersenne-schmidt - a three-mirror survey telescope," *MNRAS* **210**, pp. 597–609, Oct. 1984.

[5] Vista Web Site, "Visible and infrared survey telescope for astronomy," `http://www-star.qmw.ac.uk/~jpe/vista` , 1999.

[6] N. Kaiser, J. Tonry, and G. Luppino, "A new strategy for deep wide-field high-resolution optical imaging," *PASP* **112**, pp. 768–800, June 2000.

[7] J. Tonry, B. E. Burke, and P. L. Schechter, "The orthogonal transfer ccd," *PASP* **109**, pp. 1154–1164, Oct. 1997.

[8] T. Tengs, M. Adams, J. Pliskin, D. Safran, J. Siegel, M. Weinstein, and J. Graham, "Five-hundred life-saving interventions and their cost-effectiveness.," *Risk Analysis* **15**, pp. 369–390, 1995.