Astronomical Data Analysis Software and Systems XV **O5.2**
ASP Conference Series, Vol. XXX, 2005
C. Gabriel, C. Arviset, D. Ponz and E. Solano, eds.

The Large Synoptic Survey Telescope

Tim S. Axelrod

LSST Data Management Project Scientist, Steward Observatory, The University of Arizona, 933 N. Cherry Ave, Tucson AZ 85721

Abstract. The Large Synoptic Survey Telescope (LSST) is an 8.4 meter telescope with a 10 square degree field degree field and a 3 Gigapixel imager, planned to be on-sky in 2012. It is a dedicated all-sky survey instrument, with several complementary science missions. These include understanding dark energy through weak lensing and supernovae; exploring transients and variable objects; creating and maintaining a solar system map, with particular emphasis on potentially hazardous objects; and increasing the precision with which we understand the structure of the Milky Way. The instrument operates continuously at a rapid cadence, repetitively scanning the visible sky every few nights.

The data flow rates from LSST are larger than those from current surveys by roughly a factor of 1000: A few GB/night are typical today. LSST will deliver a few TB/night. From a computing hardware perspective, this factor of 1000 can be dealt with easily in 2012. The major issues in designing the LSST data management system arise from the fact that the number of people available to critically examine the data will not grow from current levels. This has a number of implications.

For example, every large imaging survey today is resigned to the fact that their image reduction pipelines fail at some significant rate. Many of these failures are dealt with by rerunning the reduction pipeline under human supervision, with carefully "tweaked" parameters to deal with the original problem. For LSST, this will no longer be feasible. The problem is compounded by the fact that the processing must of necessity occur on clusters with large numbers of cpu's and disk drives, and with some components connected by long-haul networks. This inevitably results in a significant rate of hardware component failures, which can easily lead to further software failures. Both hardware and software failures must be seen as a routine fact of life rather than rare exceptions to normality.

1. Introduction

The scientific motivation for the Large Synoptic Survey Telescope, hereafter LSST, was laid out in the National Academy of Sciences "Decadal Survey" report in 2000 (NAS 2000). This report identified several broad areas that will benefit from a large aperture optical telescope with a wide-field optical imager that scans the entire visible sky every few nights. These included mapping of dark matter, astrophysics of transient sources, and a more complete census of the solar system. In the last few years the science case has become more detailed, and there is better understanding of how LSST will contribute to dark matter and dark energy studies through weak lensing analysis, discovery of large numbers of supernovae, and counting of galaxy clusters. Additionally, the mapping of the structure of the Milky Way through stellar proper motions has become an important part of the science case.

To achieve these science goals will be a challenge. The LSST must observe with high efficiency, minimizing the time that is taken moving from field to field as it scans the sky. It must maintain superb image quality across an unprecedently wide field for a telescope of its size. And it must be able to handle a data stream that delivers 17 TB of raw image data every clear night, reliably processing the images to deliver astronomical data products that are freely available to the astronomy community, and to the public.

The LSST is currently in a design and development phase that is funded through 2009. Construction is expected to begin in 2009, with the system achieving first light in 2012. The design presented here is therefore preliminary.

This paper is organized as follows: The design of the LSST telescope and camera are briefly sketched in Section 2. Then the data management system (DMS) is discussed more fully in Section 3, with a particular focus on the data products that will be available to the community. The paper concludes with a discussion of the computing challenges that must be met to successfully build and operate the LSST.

2. Telescope and Camera Systems

The science case for the LSST sets the telescope etendue (the product of the area of the clear aperture and the angular area of the field on the sky) at 300 $m^2 deg^2$. This value greatly exceeds that of any current optical telescope, and is the major driver of the design. The other main requirements are for a flat focal plane, so that it is practical to use large ccd mosaics, and excellent image quality across the entire field of view and over the full range of optical wavelength. The baseline optical design (Fig 1 and 2) satisfies all the design requirements (Liang et al 2005). It has an aperture of 8.4m, a field of view of 3.5 degrees, and a complement of filters extending from u to z. The main features to note are the M1 and M3 cast together as a single piece of glass, the fast F-number, and the large aperture of the corrector elements in the camera. All of these pose fabrication challenges, but overcoming them is within the current state of the art.

The telescope itself is shown in Figure 3 (Krabbendam et al 2005). To maintain optical quality, it must actively control the figure of all three mirrors. The structure is designed to be very stiff to ensure that vibrations are damped quickly following a telescope slew. Several telescope sites are under consideration, two in Chile and one in Mexico.

The camera design is as challenging as the optics (Kahn et al 2004). Adequate sampling of the large field of view in good seeing dictates a very large number of pixels - more than three billion of them. The science case requires that the LSST move rapidly from field to field to enable the repetitive scanning of the entire visible sky every few nights, and this in turn sets a stringent limit of

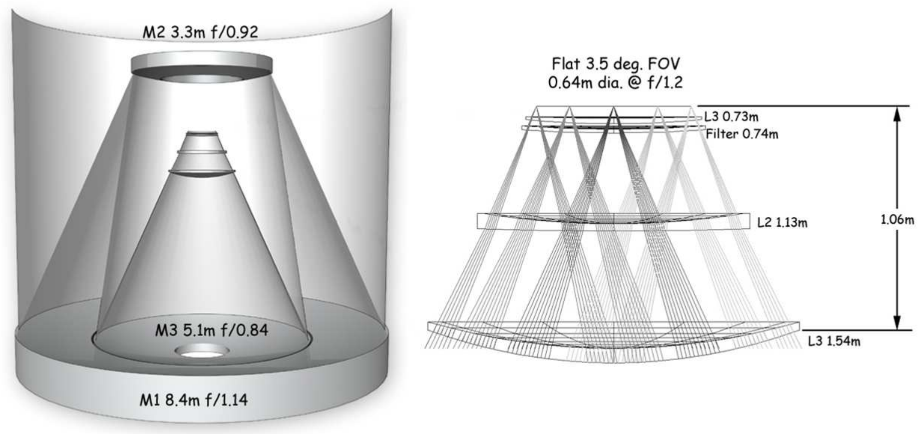

Figure 1. LSST optical design

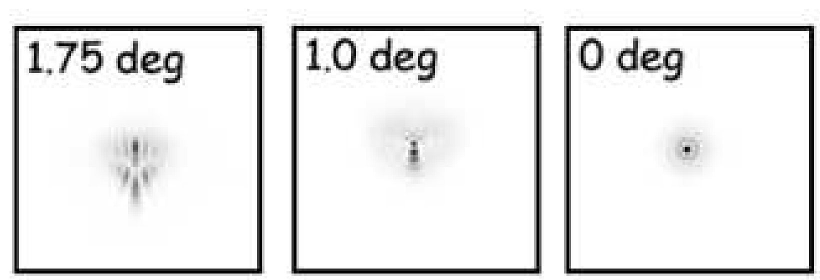

Figure 2. Point spread function - boxes are 0.6 arcsec on a side

2 seconds on the camera readout time. The baseline focal plane layout, shown in Figure 4, is based on ccds, but we are also investigating the use of cmos sensors. As shown in the figure, the focal plane meets the fast readout requirements by using a large number of parallel readout channels - 6432 of them.

3. Data Management System

The Data Management System (DMS) of the LSST, has many functions that are similar to those of previous large imaging surveys, such as MACHO, 2MASS, and SDSS. However, there is an enormous quantitative difference in the data volume. Large imaging surveys today produce a few GB per clear night. The LSST, with its 3 Gpixel imager and rapid observing cadence, produces roughly 17 TB per clear night, an increase of over 1000x. As we explain below, this quantitative increase results in some qualitatively new challenges in the data management area. Data reduction algorithms must improve because the LSST science requirements contain stressing limits on photometric and astrometric accuracy. Arguably even more challenging is the need to ensure data quality in this enormous volume of images, and in the information derived from them.

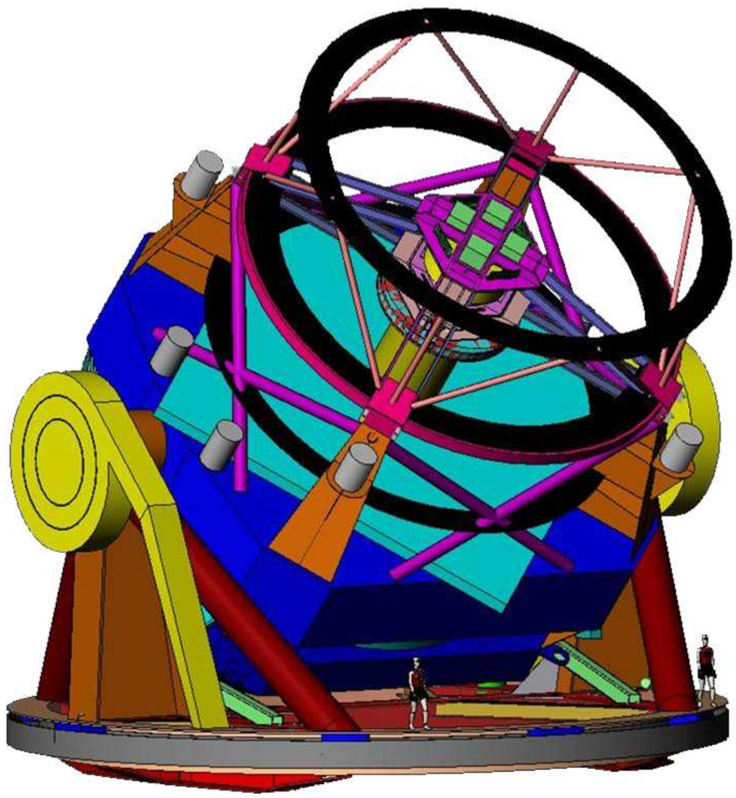

Figure 3. LSST telescope design

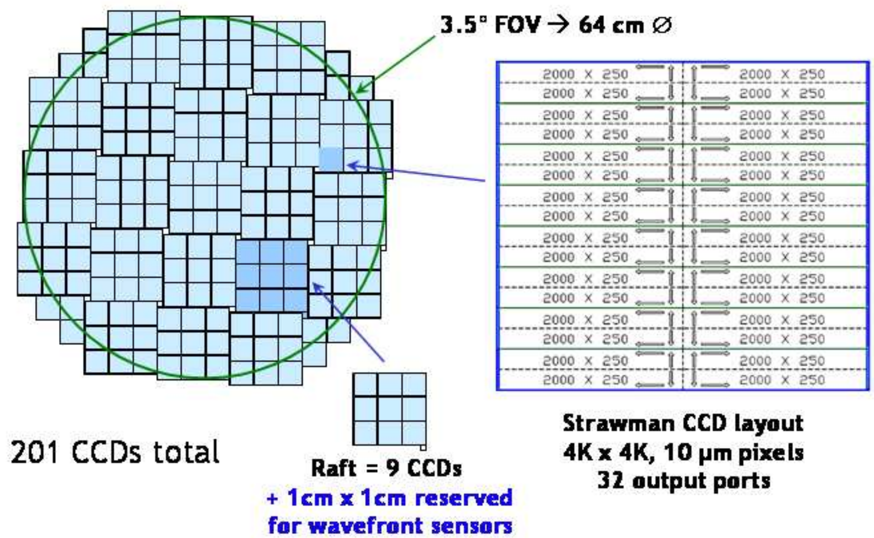

Figure 4. Layout of the LSST focalPlane

Finally, the LSST is committed to making all data freely available to the public, without imposing proprietary periods. As many previous surveys have found, this is a challenging job in itself, made considerably more difficult by the volume of LSST data. A broad overview of the DMS is available in (Kantor et al 2005).

The overall organization of the DMS hardware infrastructure is shown in Figure 5 (Dossa et al 2005). A tightly coupled mountain/base facility collects the raw camera data, and performs reductions sufficient to detect transient events that may require rapid followup by the astronomical community. The science requirement is that these reductions be complete within 60 seconds of the camera shutter closing. We are currently planning to send alerts to the community using the evolving VOEvent standard.[1]

The raw data is transferred over a high speed long haul network to an archive center, which will be located in the continental US. In addition to being the primary repository of all the survey data, the archive center performs data reduction tasks that are impractical to perform at the base facility, since they involve combining information from a large number of observations. This processing includes production of periodic major data releases, and reprocessing of the entire survey data as required to incorporate improvements in software.

Although the LSST Project is committed to producing a set of core science data products, the usefulness of the data will be extended by the larger com-

[1] `http://www.ivoa.net/twiki/bin/view/IVOA/IvoaVOEvent`

munity well beyond these core products, in ways that cannot be fully predicted. Given the enormous size of the LSST data set, producing these additional data products will require a large investment of computing resources. This is the role of the LSST Data Center depicted in Figure 5, and we assume in the architecture that there may be multiple data centers with varying capabilities. Even though these data centers must of necessity be externally funded, they will, to a greater or lesser extent, need to interface very directly to LSST DMS facilities, so it is important that we plan for them now. An overview of LSST data access plans is given in (Becla et al 2005)

A very high level view of the LSST data pipeline is shown in Figure 6. The nightly reductions involve four connected pipelines. The Image Processing Pipeline (Becker et al 2005) is responsible for producing calibrated science images, including initial astrometric and photometric calibration (these will be improved later at the Archive Center). Additionally, this pipeline is used for producing subtracted images, which play a key role in detecting transient and moving objects, and stacked images, which are used as templates in image subtraction. The Detection Pipeline (Axelrod et al 2005) is responsible for producing the Source Catalog, which contains parameters of all sources detected in an image, including location, brightness, and shape. When reducing nightly images at the base facility, the Detection Pipeline takes subtracted images as input. The same pipeline, run at the Archive Center, extracts additional information by using stacked images. The Association Pipeline (AP) (Axelrod et al 2005) is responsible for producing the Object Catalog, by associating observations of a single astronomical object from different epochs and possible different filters. The Object Catalog contains a wide variety of information about objects, including their light curves, best static estimates of shapes, proper motions, etc. The AP produces only a subset of this information when running at the base facility. The remainder is produced at a less frequent cadence at the archive center. Rapidly moving objects are processed in a separate pipeline, which ultimately produces orbital elements for the Object Catalog (Kubica et al 2005).

Two final pipelines, the Classification Pipeline, and the Deep Detection Pipeline, run only at the Archive Center. The Classification Pipeline is charged with producing the best available classification for each object in the Object Catalog. To do this, it makes use of the object's lightcurve, colors, motion (if applicable), and shape (if applicable). Additionally, it may call on external resources, such as imagery in other bands, available through the Virtual Observatory. In marginal cases, the classification result may be explicitly probabilistic, producing several possible classifications. The Deep Detection Pipeline (Roat et al 2005), is the source of most information concerning galaxies, including shape information to be used in weak lensing studies, and photometric redshifts. It is based on optimal combination of all imagery available for each galaxy.

4. Computing Challenges

We conclude with a survey of the major challenges that we face in the design and implementation of the DMS.

The LSST DMS will begin construction in 2009 and will not begin operation until 2012, but is being designed now. Computing technology will undergo

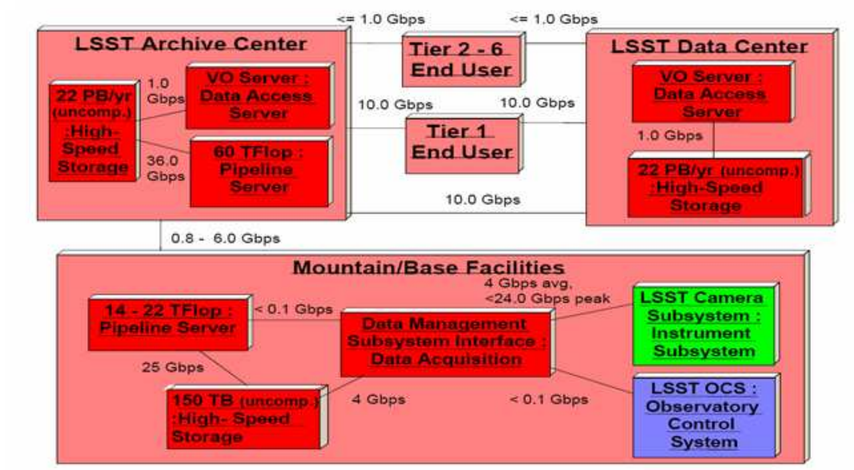

Figure 5. DMS hardware infrastructure

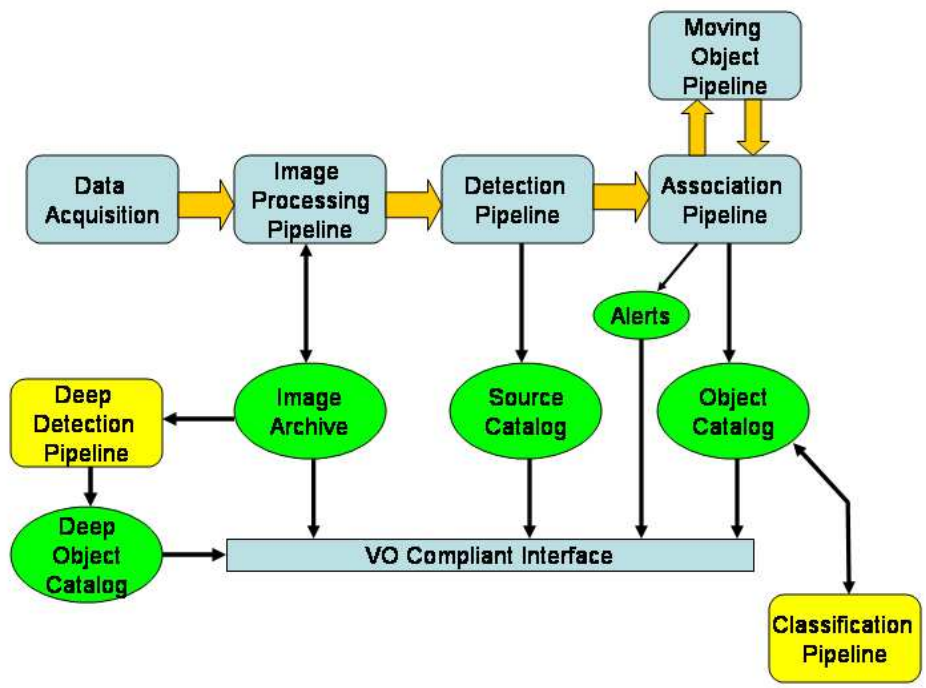

Figure 6. Top level view of the LSST DM pipeline

an enormous evolution before the LSST starts operating, and that evolution can only be imperfectly predicted. One major challenge is to account for this evolution in planning the design, being neither too conservative (which results in a system with unrealistically high predicted cost), or too optimistic (which increases risks). Our current estimates (Figure 5) are that the LSST will require an aggregate computing capability of about 80 TF, distributed between the base and archive centers. Based on computing industry predictions of computer power growth[2], the LSST in 2012 will not be a particularly large computing center, likely failing to make it onto the list of the top 500 systems. However, this level of throughput can be achieved only with a high degree of parallelism, on the order of 5000 - 10000 cpu nodes.

So a second challenge is to ensure that the LSST software provides the degree of parallelism needed to make effective use of the computing resources that will be available. At the earliest stages of the processing pipeline, this parallelism is easy to achieve. The number of independent ccd readouts is 6432, and many image processing tasks can be performed in an "embarassingly parallel" mode. This happy situation does not continue throughout the system, however. For example, the Association Pipeline will be database intensive, and will have widely varying amounts of work to perform depending on the details of the objects it is processing. How to maintain the needed parallelism through this phase is a subject of current design and development work.

These issues are not particularly new, and have been successfully dealt with by previous projects from several application domains, including computational physics and high energy physics, as well as astronomy. The most significant challenges for LSST lie elsewhere. As mentioned previously, LSST will produce data at a rate about 1000 times greater than that from previous large surveys. However, its operating budget, and therefore the number of people available to examine the data, will not grow significantly, if at all. This will force a qualitative change in the way the survey deals with its data, which will require major advances in astronomical software design.

Current pipelines for the reduction of imaging data from large mosaics have a significant failure level. Typically at least 1 percent, and often a much larger fraction, of the input images are not processed completely successfully. For example, sky levels may not be properly matched between detectors, or the determination of the world coordinate system (WCS) may fail. Best practice today is to funnel such failed reductions to a group of humans who repeat the reductions by hand, with altered parameters or algorithms, in an attempt to derive scientifically useful data from the images. The data volume of LSST will definitively rule out this approach: there will not be enough human resources to carry out the task. There are three ways in which we can deal with this situation, and probably all of them will be employed to some degree:

- Devise reduction algorithms with lower failure rates
- Automate failure recovery
- Explictly accept a well defined level of processing failure, treating it as an irreducible system inefficiency, similar to weather

[2]`http://www.top500.org`

Certainly it seems appealing to rely entirely on the third option. Surely a reduction of a few percent in system throughput is tolerable! This is undoubtedly true, but there is still a difficulty. Reduction in throughput may be tolerable, but contamination of good data with a few percent of bad data is definitely not. It is still necessary to recognize bad data with very high reliability, and in a fully automated fashion. This is likely a significant portion of the effort required to automate failure recovery, so its difficulty must not be underestimated.

More generally, some level of system failure will always be present in a system with the hardware and software complexity of the LSST DMS. Not only algorithm failures, as noted above, but a variety of hardware failures, such as failed cpu nodes and disk drives, will be a constant fact of life. In designing the system, we cannot afford to treat such failures as isolated anomalies that will be dealt with on an ad hoc basis when they occur. Again, these issues are unique neither to LSST nor to the domain of astronomy.

In fact, there are major initiatives underway in the computing industry to solve some of the more ubiquitous issues. Many of these initiatives fall under the umbrella of "autonomic computing", a phrase coined by Paul Horn of IBM in 2001[3]. There are now a wide variety of products intended to make systems that self-configure, self-heal, self-optimize, and self-protect in the face of a variety of failures and a changing external environment, such as network weather. Most of these new capabilities are aimed at the hardware or operating system level, rather than the application or algorithm level. But, they are likely to be important building blocks for the LSST DMS. Some of the techniques they employ may end up being applicable at higher levels of the system as well.

References

Axelrod, T. et al, 2005, BAAS, 207, 2630

Becker, A. et al, 2005, BAAS, 207, 2629

Becla, J. et al, 2005, BAAS, 207, 2633

Dossa, D. et al, 2005, BAAS, 207, 2634

Kahn, S. et al, 2004, BAAS, 205, 10806

Krabbendam, V. et al, 2005, BAAS, 207, 2622

Kubica, J. et al, 2005, BAAS, 207, 2632

Liang, M. et al, 2005, BAAS, 207, 2620

National Academy of Sciences 2000, "Astronomy and Astrophysics in the New Millenium"

Roat, C. et al, 2005, BAAS, 207, 2631

[3] `http://www.research.ibm.com/autonomic/manifesto/autonomic_computing.pdf`