Fermi National Accelerator Laboratory

FERMILAB-Conf-99/178-A

The Sloan Digital Sky Survey Date Acquisition System, and Early Results

J. Annis et al.

Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510

June 1999

Published Proceedings of the *11th IEEE NPSS Real Time Conference,*
Santa Fe, New Mexico, June 14-18, 1999

Operated by Universities Research Association Inc. under Contract No. DE-AC02-76CH03000 with the United States Department of Energy

The Sloan Digital Sky Survey Data Acquisition System, and Early Results

J. Annis, J. Bakken, D. Holmgren, D. Petravick, R. Rechenmacher
Fermilab, P.O. Box 500, Batavia, Illinois 60510

Abstract

The Sloan Digital Sky Survey will systematically map one-quarter of the sky, producing detailed images in five color bands and determining the positions and absolute brightnesses of more than 100 million celestial objects. It will also measure the redshifts of a million selected galaxies and of 100,000 quasars, yielding a three-dimensional map of the universe through a volume one hundred times larger than that explored to date. The SDSS collaboration is currently in the process of commissioning the 2.5-meter survey telescope. We describe the data acquisition system used to record the survey data. This system consists of twelve single board computers and their associated interfaces to the camera and spectrograph CCD electronics, to tape drives, and to online video displays, distributed among several VME crates. A central UNIX computer connected to the VME crates via a vertical bus adapter coordinates the system and provides the interface to telescope operations. We briefly discuss results from the observing runs to date and plans for the archiving and distribution of data.

I. Introduction

The Sloan Digital Sky Survey (SDSS) will image π steradians about the north galactic cap in five filters and acquire one million spectra using dedicated telescopes at the Apache Point Observatory in New Mexico. We describe the data acquisition system for the survey's three main instruments: the 54-ccd imaging camera [1], the 660-fiber twin spectrographs which use 4 ccds, and the single-ccd photometric calibration camera.

II. Instruments

The SDSS uses facilities and two dedicated telescopes at the Apache Point Observatory near Alamogordo, New Mexico. Imaging and spectroscopic data are obtained on the SDSS 2.5m telescope, which images an area on the sky equivalent to about 30 full moons (3 degrees) onto a 1-m focal plane. The SDSS 0.5m photometric calibration telescope continuously measures standard stars during observations in order to determine photometric calibrations.

The spectrographs and photometric calibration camera do not impose any extraordinary system requirements. On the other hand, the SDSS imaging camera is an extraordinary instrument, employing 54 ccds and producing data at an aggregate rate of 9 Mbyte/second. The imaging camera drives the data acquisition requirements and is shown in figure 1. The camera contains a photometric array consisting of 30 2048 x 2048 ccds arranged in 6 columns of 5 detectors each. Photometric filters of five different passbands are attached above each ccd to a corrector element; thus, the camera simultaneously images the sky in five colors. The remaining 24 ccds make up the camera's astrometric array. Each is 400 x 2048 pixels in size and is covered with a neutral density filter. These chips will not saturate when imaging bright astrometric standard stars, whose positions are well known. Two of the astrometric arrays ccds serve as focus chips; these are mounted about 200 microns behind the focal plane, and are half-covered with a window of 400 microns optical path thickness. A comparison of images between the two halves will yield a differential measurement of the focus.

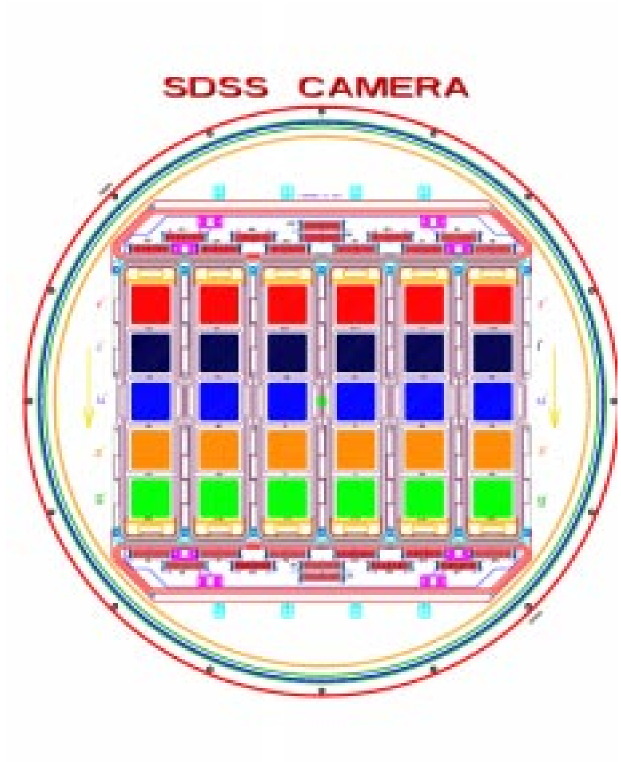

Figure 1. SDSS Imaging Camera

The imaging camera is controlled over a serial line and generates its own internal timing. The ccds are operated in Time Delayed Integration mode (TDI). The tracking of the 2.5m telescope is controlled such that, combined with the effect of the rotation of the Earth, the image of the sky moves vertically down the rows of the ccd detectors at precisely the rate at which ccd scan lines are readout. After each scan line is converted by the camera electronics, the pixel data are are transmitted to the data acquisition system over 10 fibers: one fiber per column of photometric ccds, and one fiber per row of astrometric ccds.

Imaging data are recorded in their entirety and are

Work supported by the U.S. Department of Energy under contract No. DE-AC02-76CH03000.

blocked into frames of 1354 rows from each ccd. This number corresponds to half the distance in rows between ccds in a scan line, which allows the first frame from the second chip in a column to contain the same part of the sky as the third frame from the first chip, and so forth. The frames are written to tape such that all data from a single photometric ccd column are on the same tape, the frames from corresponding parts of the sky are written together, and the frames are written in (nearly) standard FITS format.

III. Hardware Architecture

The SDSS data acquisition system is partitioned into online systems and host systems. The host systems are the root of system control and provide the user interface, while the online systems handle real time high rate data flow. The online computers handle the complete data stream from the instruments, compute summary data, and serve subsets of the data to the host computer. Their computing resources are carefully matched to the problem at hand, and their software is not subject to short timescale changes by the astronomers, unlike the software on the host machines.

The online systems are built around a number of VME backplanes connected to each other and the host system by a VME interconnect. Disk and tape are connected to separate SCSI buses. The systems run the VxWorks operating system. The current host systems are Silicon Graphics Incorporated Crimsons, each with 150 MHz R4400 processors and 144 MB of system memory. One host system controls the photometric, astrometric, and spectroscopic online systems, and the other controls the monitor telescope online system.

The data path is as follows: the camera produces a pixel value. It is transmitted, along with up to 12 other pixels from other amplifiers and other ccds, via a TAXI/FOXI optical fiber system to the online system. There, a Fermilab-built VCI+ module containing a FOXI receiver collects the pixel data into buffer memory to collate the pixel stream into single lines from single ccds. The VMEbus allows the MVME167 single board computer to access the line data, perform computations, and write the lines into frame files on the DA pool disk. Simultaneously, the MVME167 drives a Vigra MMI 250 VMEbus graphics controller to display the lines on video monitors ("scrolling displays") as the lines are completed. Once frames are completed the MVME167 writes them redundantly to twin digital linear tape (Quantum DLT2000) drives. The observers have access for more detailed analysis to the completed frames over the VMEbus High Performance Data Network (PT-VME940, Performance Technologies, Inc.), which supports 30M byte/sec transfers to the host computer.

The imaging camera data acquisition subsystem consists of three VMEbus backplanes, two photometric and one astrometric. Each photometric crate (figure 2

has three MVME167s (33 MHz, 32 MB system memory), VCI+ modules, Vigra boards, and Ciprico "Rimfire" SCSI controllers, and has an associated disk pool and 6 DLT tape drive farm. The astrometric crate has four MVME167s,

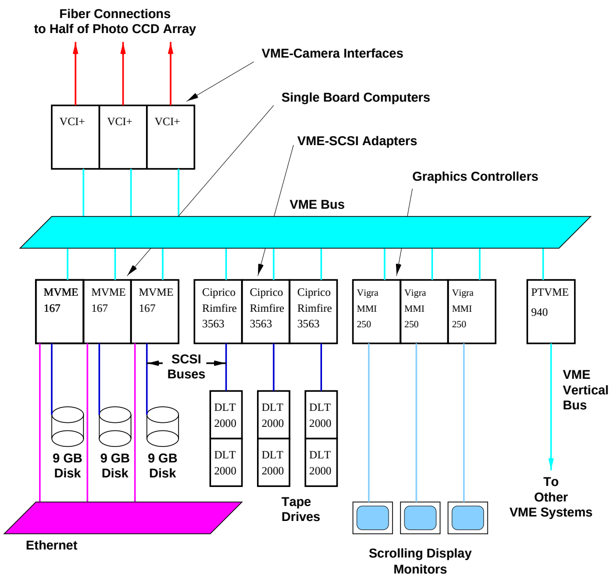

Figure 2. Photometric VMEbus Layout

VCI+ modules, Vigra boards, Vigra boards, and Rimfire SCSI controllers and an associated disk pool. The spectroscopic DA consists of a single VMEbus backplanes MVME167, Vigra board, and associated disk pool, but has two VCI+ modules to accept non-synchronous data from the twin spectrographs. The photometric calibration camera DA has a single VMEbus backplane with a MVME167, VCI+ module, Vigra board, and associated disk pool. Images from the latter two systems are sent to the host computer via the PT-VME940 link and are written to DLT2000 tape there.

IV. Software

Nearly all of the survey's software is built around the Tool Command Language (TCL) developed at the University of California, Berkeley, by John Ousterhaut et al.[3] TCL is a C-based user extensible command interpreter. Survey command primitives are written in C, and declared to a TCL interpreter. Observing programs can be constructed in TCL from these primitives. The fundamental off-line toolkit for SDSS is Dervish. It provides the C and TCL frameworks to access, manipulate, analyze, and display *regions*, which may be for example frames from the imaging camera, or 2D spectra from the spectrograph. Dervish integrates the *libfits* FITS manipulation library, plus Fermilab versions of *pgplot*, *saoimage*, and *TCL*. Off-line packages built upon *Dervish* include *astrotools*[2], *sdssmath*, and *fslalib*, which together provide common standardized numerical routines and coordinate conversions and other common tools for astronomers. These in turn are based upon *Slatec*, *Lapack*, and *FFTW*.

The data acquisition software package, *Astroda*, simultaneously runs on the VxWorks online and IRIX host systems. *Astroda* is written in ANSI C. Its architecture includes multiple TCL interpreters on the online computers and regions of memory shared between the host and online systems via the PT-VME940 link.

The activities of *Astroda* on the online computers are organized into tasks, which communicate to each other using messages and status entries.

- The *line* task is woken by the VCI+ interrupts as each data line for the ccds is available. It rearranges the received pixels into frame buffers for each ccd.
- The *assem* task writes complete frame buffers to the pool disk, thus assembling frames.
- The *astroline* task histograms all columns; finds and cuts out patches around bright stars; as time permits calculates quartiles, bright star parameters, and sky values; assembles the cut-out stars (*postage stamps*) and derived quantities into FITS files called *gangs*; and computes focus measures.
- The *archiver* task notes when frames are completed and programs the Rimfire SCSI interfaces to queue the frames to tape in field order.
- The *data server* task listens for requests for frames and gangs from the host computer and programs the PT-VME940 board to deposit the data into host system memory.
- The *command server* task listens for general control messages.
- The *scrolling display* task advances the scrolling display as each line is received, first depth reducing the data from 16 to 8 bits.

Each MVME167 maintains a status database in its local memory that is, importantly, accessible from the host system. A status entry is identified by an alphanumeric name, a type (integer, float, etc.), the actual, minimum, and maximum values, a description, protection and current validity. There are about 235 status entries.

The online systems report error and status to the host computer over an Ethernet connection using the Fermilab Murmur package.[4]

The host system runs the *astroda* program in order to communicate with the online systems and obtain access to the data over the PT-VME940 board.

Observing programs, programmed by SDSS astronomers, are based on *astroda* and run on the host systems. They must operate complicated instruments, monitor instrument health, and ensure high data quality. The imaging observing program, or *IOP*, for example, monitors over 10,000 quantities including the states of the camera microprocessors and liquid nitrogen fill systems, the voltages on all of the instruments, the status of the online tape archiver and disk pool, the status of *astroline* analysis, the status of the telescope and interlock systems, and the weather. All monitored quantities are available to an observer as trend plots. If any of these monitored values are out of tolerance, or any microprocessor reports a state it believes to be bad (e.g., the dewar temperature rising), *IOP* denotes this by a hierarchy of red, leading down to the subsystem reporting a problem, and topping out in a set of 12 very high level red/green lights.

V. Lessons Learned

Even the most careful schedules prove to be inaccurate on a project the size of SDSS. The data acquisition hardware and software were completed and delivered to the observatory in 1996, about 20 months before first light on the 2.5m telescope (May 1998). In anticipation that at least parts of the DA would be ready before the ccd camera, camera simulators were built which have the ability to send data stored on disk to the VCI+ boards via interfaces identical to the real cameras. Through the use of these simulators, we were able to integrate the data acquisition software, including mock observers programs, and extensively test both the DA hardware and software. Prior to real imaging data, SDSS astronomers provided simulated frames which were used to test and refine *astroline*. Following first light, the simulator environment was used to re-run real data through the system in daytime to allow further refinements of *astroline*.

End-to-end error checking also proved to be of considerable value. *Astroda* codes include optional checksumming at various critical points in the data flow (for example, before writing to disk, after reading from disk, after reading from tape). Comparing simulated data entered into the system bit for bit with data read from DLT tapes revealed various software and hardware bugs, and gave us great confidence in the robustness and reliability of the system. This error checking revealed read and write errors on an earlier brand of pool disk which were shown to be disk firmware errors.

VI. Early Results

The first images were obtained on the 2.5m telescope on May 27, 1998. We are currently in the survey's test year, devoted to commissioning the various instruments and bringing the systems to a state of operational readiness for the formal start of the five year survey in 2000. In imaging data from very early commissioning runs, 15 high-redshift quasars were identified,[5] four with among the highest redshifts observed ($z = 4.63, 4.75, 4.90, 5.00$). Two methane brown dwarfs have been identified; only one other such object has previously been observed. A faint high latitude carbon star was identified, confirming that SDSS will efficiently find large numbers of FHLCs. The first spectrum and redshift were obtained on May 25, 1999.

VII. Data Archives and Distribution

SDSS has two databases, both implemented with the Objectivity OODBMS.[6] The operations database is maintained at Fermilab, tracking the state of the survey, the observing schedule, the objects detected during specific runs, and so forth. The science database (SDBX), maintained at Johns Hopkins University, contains the calibrated object information.

The DLT tapes containing the raw data from a night's observing are shipped daily to Fermilab via an express delivery service. Analysis pipelines are run against these data at Fermilab, which calibrated parameters shipped via the internet to the SDBX at JHU. *Target selection* from data in the SDBX identifies objects which will be assigned fibers for subsequent spectroscopic observation.

The volume of data to be obtained by the survey is modest by high energy physics standards, but substantial nonetheless. The imaging data will total 8.2 Terabytes, consisting of a 0.4 arcsecond resolution pixel map of the sky in five colors. The spectra data will total 50 Gbytes. Altogether the survey will identify and parameterize 1 billion objects, including 100 million galaxies and 1 million quasars, with spectra measured for 1 million selected objects.

Data will be available according to a published schedule.[7] Initially this will occur in batches of calibrated photometry and spectroscopy every one and a half years, with the first release expected in mid-2001. Towards the end of the survey additional data will be released yearly. The various data products, their sizes, and their release form are shown in Table refdata.

TABLE I
SDSS Data Products

Product	Size	Form
Redshift Catalog	2 GB	CD-ROM
Survey Description (Status, Calibrations)	1 GB	CD-ROM
Compact Photometric Catalog	60 GB	CD-ROM
1D Spectra	60 GB	On-line
Atlas Images	1.5 TB	On-line
Compressed Sky Map	0.3 TB	On-line
Full photometric catalog	400 GB	On-line

VIII. References

[1] Gunn, J.E., et al, "The Sloan Digital Sky Survey Photometric Camera," *The Astronomical Journal*, vol. 116, pp.3040-3081, December 1998.

[2] Kent, S., et al, "Sloan Digital Sky Survey," *Proceedings of the Third Annual Conference on Astronomical Data Analysis Software and Systems.*

[3] Ousterhaut, J., "Tcl: An Embeddable Command Language," *Proceedings of the Winter 1990 USENIX Conference.*

[4] Oleynik, G., et al, "Murmur - A Message Generator and Reporter for UNIX, VMS, and VxWorks," *IEEE Real Time '93 Conference Record.*

[5] Fan, X., et al, "High-Redshift Quasars Found in Sloan Digital Sky Survey Commissioning Data," *The Astronomical Journal*, in press, July 1999.

[6] Objectivity Incorportated, 301B East Evelyn Avenue, Mountain View, California, 94041.

[7] `http://www.sdss.org/science/schedule.html`.