Sloan Digital Sky Survey
Quarterly Progress Report
Fourth Quarter 2004

February 8, 2005

Table of Contents

Q4 PERFORMANCE HIGHLIGHTS

- We revised the Five-Year Baseline Plan by recasting imaging progress on the North Galactic Cap in terms of “footprint” square degrees. The original baseline plan tracked imaging progress in terms of “unique” square degrees. The change was necessary to reveal the remaining imaging area between stripes 10 and 37, and is described in detail in Section 1.1.
- We obtained 1,071 square degrees of new imaging data against a baseline goal of 1,160 square degrees (92%).
- We completed 31 standard survey plates, six special program plates, and eighteen SEGUE plates. Converting these to plate-equivalents, we completed 75 plate-equivalents against our baseline goal of 153 plates (49%).
- Weather was marginal in October (51%), poor in November (34%), and marginal in December (56%). The shortage of suitable weather conditions impacted our observing yield.
- Operations continue to run smoothly. There were no major system problems at APO. All data acquired during the quarter, including SEGUE and Supernova program data, have been processed. Plug plates continue to be designed, fabricated, and delivered to APO in a timely fashion.
- The initial version of DR4 was made available to the SDSS collaboration on December 22. DR4 contains 6,670 square degrees of imaging data (180 million objects), compared to 5,282 square degrees of imaging data (141 million objects) in DR3. DR4 also contains 672,649 spectra (including 478,890 galaxies), compared to 528,640 spectra (including 374,730 galaxies) in DR3. The public release of DR4 is scheduled for July 2005, in accordance with the approved data distribution plan.
- Data archive use continues to grow. Through the end of December, the public SkyServer had received just over 65 million web hits and executed 9.3 million SQL database queries. We averaged 2.9 million web hits per month in Q4, compared to 2.7 million hits in Q3 and 2.2 million in Q2. We averaged 642,000 SQL queries per month in Q4, compared to 614,000 queries in Q3 and 577,000 in Q2. For the twelve months ending December 31, 2004, we achieved 99.5% uptime on our public database cluster.
- Q4 cash operating expenses were $690K against a baseline budget of $767K (-10%), excluding management reserve. In-kind contributions were $299K against anticipated contributions of $450K (-34%). No management reserve funds were expended in Q4. Unspent management reserve funds will be carried forward once final 2004 invoices have been settled.

1. SURVEY PROGESS

1.1. Revision of the Five-Year Baseline Plan

The SDSS Five-Year Baseline Plan provides the expected rate of data collection with respect to time, separately for imaging and for spectroscopy, and separately for the North Galactic Cap region and for the fall, or "Southern," sky. It was developed in 2001 to provide a way to measure actual progress on the survey with respect to reasonable goals. A benchmark is useful to the extent it does not shift around; we have been careful to maintain the 2001 plan for this reason, despite significant changes in strategy (for example, undertaking imaging in the final observing season, whereas the Plan assumed that the final season would consist of only spectroscopy).

Until now the Five-Year Baseline Plan has served its purpose well, providing a consistent basis for reporting in our quarterly reports and annual reports. However, a feature of the accounting for imaging data, which was originally sensible, is now leading to an inaccurate picture for the status of the survey. (These statements do not apply to the spectroscopic survey or to the imaging and spectroscopy in the South Galactic Cap.) We therefore propose to re-cast the Five-Year Baseline Plan for imaging in a way that properly shows the status of the sky coverage, leading up to the completion of the SDSS. We emphasize that the survey goal is not being relaxed: the change concerns only the way in which we account for the imaging data.

The original accounting used a category we call "unique" imaging data. This category corrects for overlaps at the ends of strips and in between strips in a stripe. Each hour of observing time results in (potentially) 18.75 square degrees of "unique" imaging data. The direct scaling of time into area for this category was useful because all projections were based on the amount of observing time that was expected to be available.

Another category of imaging data is called "footprint," which corrects for overlaps between stripes due to the convergence of the system of great circles. We did not use this at the beginning since it does not scale in a simple way with time - each new observation will yield a different number of "footprint" square degrees, depending on where it falls with respect to earlier observations. (As a rough indication, "footprint" area is smaller than "unique" area by about 20%.) At this time in the survey, we are trying to finish observations within a specific area of the sky, and the "footprint" metric is far more meaningful.

The Five-Year Baseline Plan was a projection of the amount of data that could be collected within the North Galactic Cap, but it made no explicit assumptions of where the data would be within the pi steradians of the NGC. We are motivated to leave a legacy of a survey within a contiguous area, specifically completing stripes 10 through 37. This goal is more demanding: we now have a goal not just of a total area, but of a particular area. Nevertheless our progress on the imaging survey is very close to what we had projected in the Five-Year Baseline Plan.

The need for a revision in our accounting and reporting is clear from the following: the Five-Year Baseline Plan stipulates 7700 square degrees of "unique" imaging data in the North Galactic Cap. As of 31 December 2004, we have collected 7831 square degrees of "unique" imaging data. Despite this, as of the end of 2004 Q4, we were still missing over 1300 square degrees of "unique" imaging data in the critical region between stripes 10 and 37. The new reporting is intended to reveal this remaining amount of imaging and to provide a way to track progress in completing it from now until the end of the SDSS (30 June 2005).

As just mentioned, the specific goal is to complete, in the imaging survey, the region between and including stripes 10 and 37 by 30 June 2005. This is an area of 7000 footprint square degrees. In addition, we have already scanned about 700 footprint square degrees in the NGC that are on adjacent stripes, or in the Spitzer First Look Survey region. Thus the new reporting needs to reflect a goal of having acquired a total of 7700 footprint square degrees by 30 June 2005. For reference, 7,700 footprint square degrees corresponds to approximately 9370 unique square degrees.

The simplest projection is to assume an average rate of acquiring data that sums to 7700, given that the start of the survey is defined to be 1 April 2000 (as we did before). Figure 1.1 shows this projection, and provides the actual data-collection rate in the same terms. As of the end of Q4 2004, the actual amount of "footprint" imaging data in the NGC was 6458 square degrees, which is very close to that required if we are to meet our goal in the next two quarters.

1.2. Q4 Imaging

We collected imaging and spectroscopic data for the SEGUE and Supernova programs, and on the Northern Galactic Cap as it became available. This section summarizes progress on the standard survey. Progress reports on the SEGUE and Supernova programs are presented in Section 8.

Table 1.1 compares the imaging data obtained against the revised baseline projection for the North Galactic Cap and against the established baseline projection for the South and Southern Equatorial Stripe. We collected a total of 1,071 square degrees of imaging data, or 92% of our baseline goal of 1,160 square degrees. In addition to standard survey scans, imaging including several SEGUE scans and scans on strip 82N (one-half of the Southern Equatorial Stripe) for the Supernova program.

Table 1.1. Imaging Survey Progress in Q4-2004

	Imaging Area Obtained (in Square Degrees)			
	Q4-2004		Cumulative through Q4	
	Baseline	Actual	Baseline	Actual
Northern Survey[1]	292	279	6538	6458
Southern Survey[1]	0	0	745	738
Southern Equatorial Stripe[2]	868	792	4808	3470

1. "Footprint" area
2. "Good minus Unique" area.

The following graphs show progress against the imaging goals for the Northern Galactic Cap and the Southern Equatorial Stripe. The full set of graphs is available on the SDSS website.

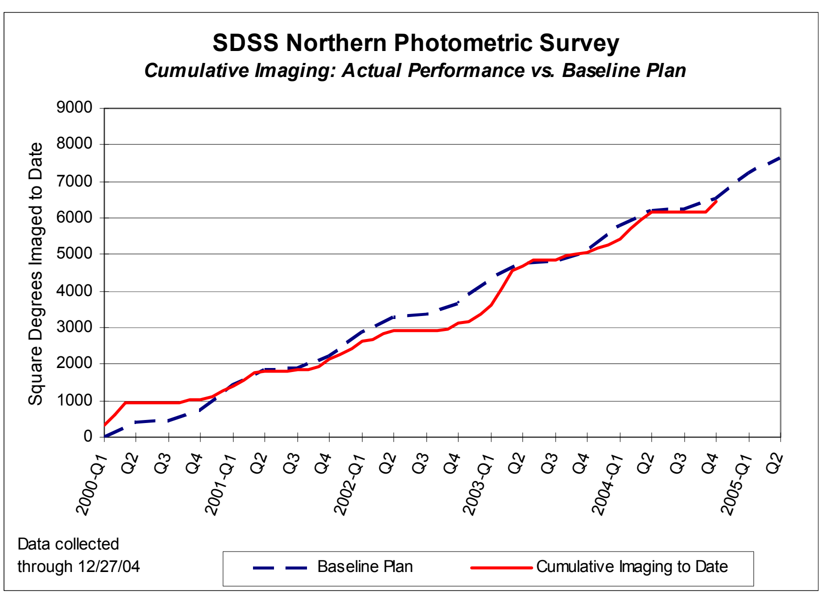

Figure 1.1. Imaging Progress against the Baseline Plan – Northern Survey

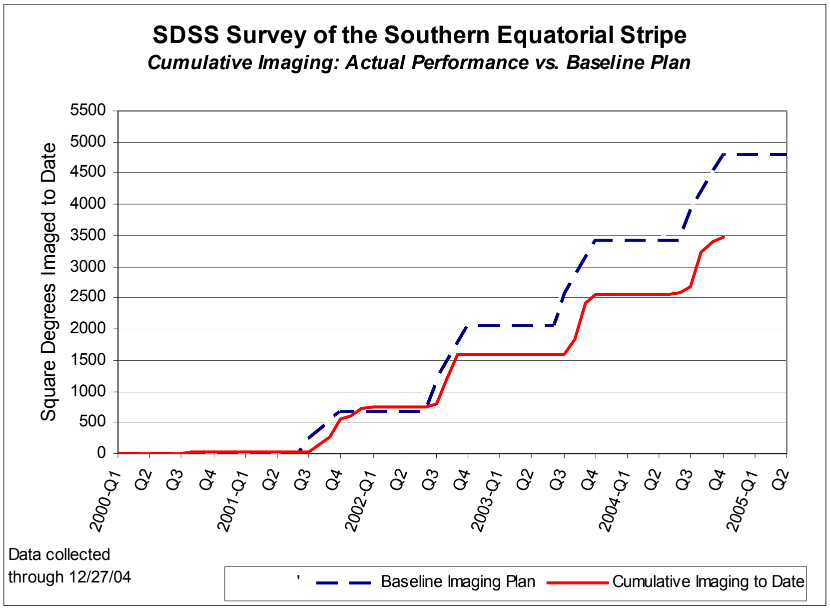

Figure 1.2. Imaging Progress against the Baseline Plan – Southern Equatorial Survey

1.3. Q4 Spectroscopy

Spectroscopic progress is reported in terms of the number of plates observed and declared done during a quarter. Each plate typically yields 640 unique spectra.

In Q4, we completed a total of 55 spectroscopic plates. Of these, 31 were standard survey plates, six were special program plates, and eighteen were SEGUE plates. Some of the special plates require longer exposure times than the standard plates, so it is necessary to convert the number of physical plates observed into plate-equivalents to measure progress against the baseline. In these terms, we completed 75 plate-equivalents against the baseline goal of 153 plates (49%).

The following graphs show spectroscopic progress against the baseline goal for each survey region. Progress is reported in plate-equivalents for the special program plates, which allows for a direct comparison with baseline goals. Details on the SEGUE plates are provided in Section 8.

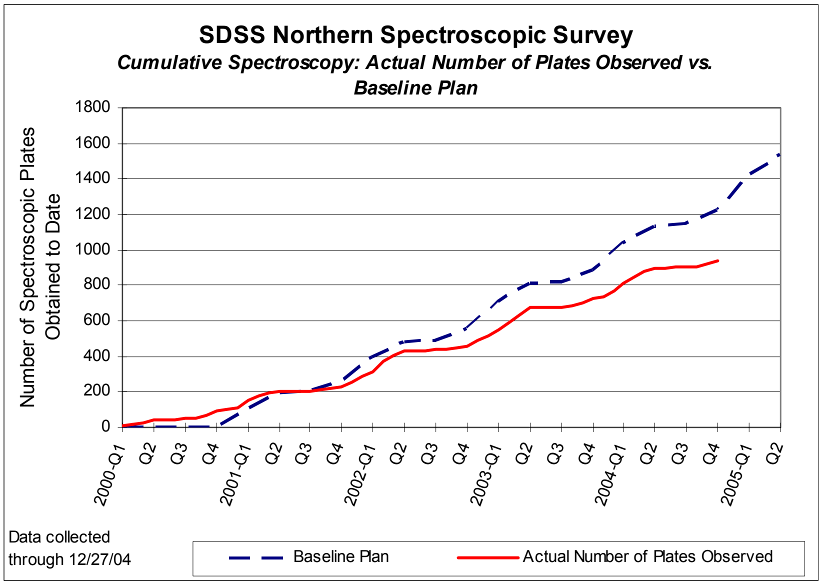

Figure 1.3. Spectroscopic Progress against the Baseline Plan – Northern Survey

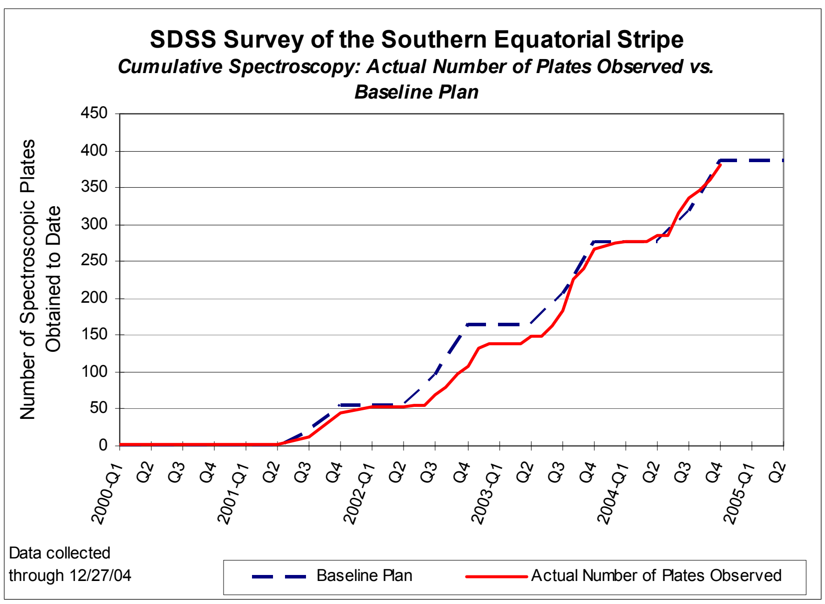

Figure 1.4. Spectroscopic Progress against the Baseline Plan – Southern Equatorial Survey

2.0 OBSERVING EFFICIENCY

The following sections summarize observing efficiency according to the categories used to prepare the baseline projection.

2.1. Weather

The weather category reports the fraction of scheduled observing time that weather conditions were suitable for observing. Table 2.2 summarizes the total amount of time lost to poor weather in Q4. Figure 2.2 plots the fraction of suitable observing time against the baseline forecast. In Q4, we lost more time to weather than predicted in the baseline plan.

Table 2.1. Potential Observing Hours Lost to Weather in Q4

Observing Condition	Total hours potentially available for observing	Total hours lost to weather	Fraction of time suitable for observing	Baseline Forecast
Dark Time	431	206	52%	60%
Dark & Gray Time	651	326	50%	---

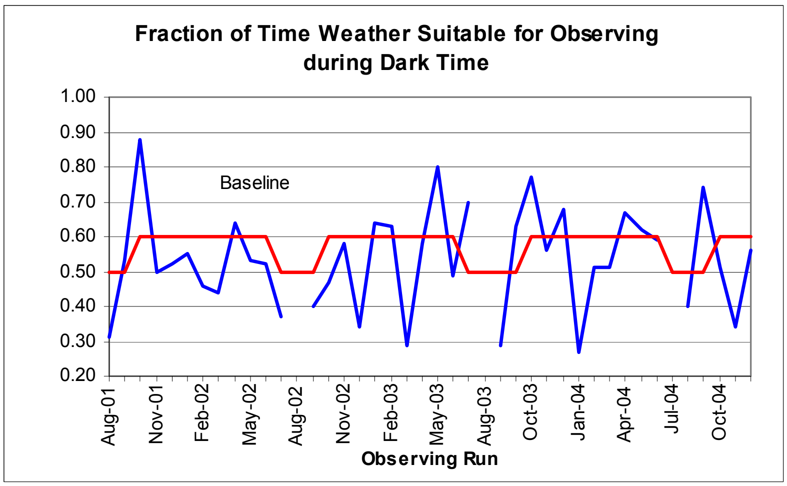

Figure 2.1. Fraction of Time Weather Suitable for Observing

2.2. System Uptime

System uptime measures the availability of equipment when conditions are suitable for observing. We averaged 92% uptime in Q4 against a baseline goal of 90%. Table 2.2 summarizes the total amount of time lost to equipment or system problems and Figure 2.2 plots uptime against the baseline over time. We lost observing time in Q4 to spectrograph collimator actuator problems, PTVME link failures, and a defective azimuth drive slip encoder. In addition, we experienced a number of smaller problems that caused difficulty during observing operations but did not result in lost observing time.

Table 2.2. Potential Observing Hours Lost to Problems in Q3

Observing Condition	Total hours potentially available for observing	Total hours lost to problems	System Uptime	Baseline Forecast
Dark Time	431	32	92%	90%
Dark & Gray Time	651	35	95%	---

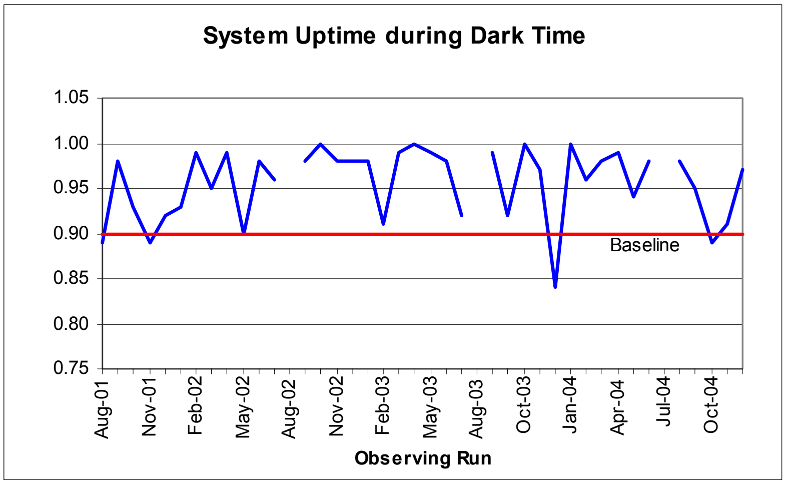

Figure 2.2. System Uptime

2.3. Imaging Efficiency

The imaging efficiency ratio provides a measure of observing efficiency and is defined as the ratio of science imaging time to the sum of science imaging time plus imaging setup time. The baseline plan established the imaging efficiency ratio to be 0.86. When the weather is poor, as it was in much of Q4, imaging efficiency drops because the observers may take more time than usual to finish setup and calibration; setup and calibration can be performed while waiting for the weather to clear. November provides a more accurate representation of our current level of imaging efficiency. In November, the observers achieved an imaging efficiency of 0.91, which exceeded the baseline target of 0.86.

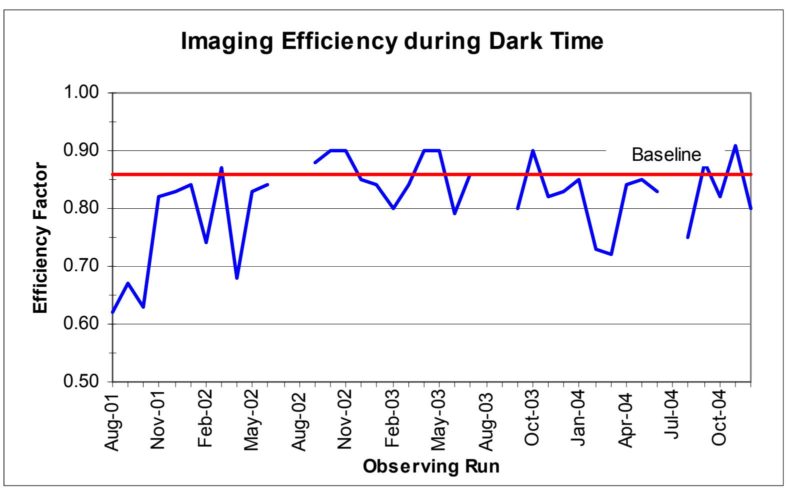

Figure 2.3. Imaging Efficiency

2.4. Spectroscopic Efficiency

Spectroscopic efficiency is derived by assessing the time spent performing various activities associated with spectroscopic operations. Table 2.3 provides the median time, by dark run, for various overhead activities associated with spectroscopic operations. Units for all categories are minutes except for efficiency, which is given as the ratio of baseline science exposure time (45 minutes) to total time required per plate. Using these measures, we averaged 64% in Q4, which matched the baseline goal of 64%.

Table 2.3. Median Time for Spectroscopic Observing Activities

Category	*Baseline*	*Run starting Oct 7*	*Run starting Nov 1*	*Run starting Dec 1*
Instrument change	10	5	5	5
Setup	10	12	10	12
Calibration	5	6	6	6
CCD readout	0	3	3	3
Total overhead	25	26	24	26
Science exposure (assumed)	45	45	45	45
Total time per plate	70	71	69	71
Efficiency	0.64	0.63	0.65	0.63

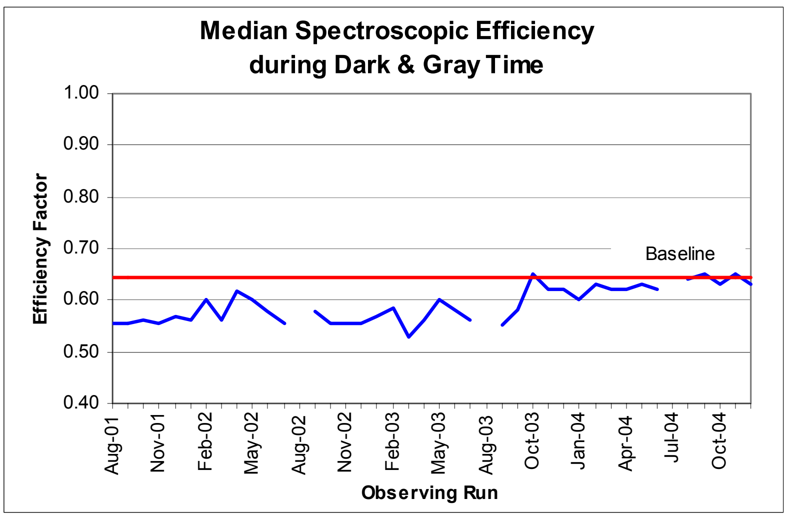

Figure 2.4. Spectroscopic Efficiency

3. OBSERVING SYSTEMS

In addition to addressing the problems described earlier, we addressed a few minor issues and completed a number of planned engineering tasks and preventive maintenance activities during the summer shutdown.

3.1. The Instruments

The imaging camera worked relatively well throughout the quarter. We did see banding in bias levels across the CCDs in dewar 1 early in the quarter. The exact cause of the banding is still under investigation; fortunately, data quality has not been negatively affected.

We saw an increase in the number of reported spikes in the imager camera vacuum in Q4. We do not believe that the vacuum is changing because of the speed at which the vacuum reportedly changes. Rather, we suspect that the ion pump or its controller may be acting up. To help identify the source of the problem, we installed a data logger on the ion pump current input. If the problem turns out to be the ion pump, we will need to go into the camera to fix the problem. We would like to postpone this to the summer shutdown if at all possible.

A series of calibration measurements were made on the imaging camera. Linearity measurements were made and found to be consistent with measurements taken in 1999. The wavelength of the imager band responses was measured and compared to measurements made in July to determine temperature sensitivity. A preliminary review of the data indicates that band widths in *i*, and possibly *r,* are approximately 2 nm broader than similar data taken in July. A more detailed report will be issued soon.

We had mechanical problems with the spectrographs in the quarter. Specifically, we had problems with the actuator assemblies for the spectrograph collimator motors, which affected spectrograph focus and impacted observing operations. We looked into replacing the actuator assemblies but found that components are obsolete and no longer available. Modifying the system to accept new actuator assemblies will be time-consuming and costly, and so is being postponed pending the approval of SDSS-II. In the interim, the engineering crew designed and fabricated tooling, and developed and implemented preventive maintenance procedures, to allow us to disassemble, inspect, and re-lubricate the existing actuator motors and gear heads. In Q4, the actuator assemblies were disassembled and re-lubed; the system has worked well since this maintenance work was performed.

3.2. The 2.5m Telescope

We performed a number of engineering, maintenance and repair tasks in Q4. The following list summarizes some of the more significant activities:

1) We have had a recurring problem with refrigerant leakage in the Thermal King unit on the telescope enclosure. We contracted with a maintenance crew from Thermal King to replace the seals on the compressor. To prevent the seals from drying out in the future, they suggested that we operate the unit daily. We subsequently installed a timer switching system that now automatically runs the unit for one hour per day. Recall that the Thermal King unit is needed to keep the telescope and mirrors close to ambient temperature when the telescope is opened daily.

2) We had a problem with one of the slip detection encoders and in diagnosing the cause, discovered corrosion on the electronic circuitry inside the housings of some of the slip detection encoders. Although the vendor did not acknowledge a corrosion problem with their units, they did note that the encoders had been re-designed. We replaced our existing encoders with the new design to preclude further problems associated with the corrosion.

3) To increase operating efficiency during imaging operations, we extended the normal operating range of the telescope azimuth axis from +/-270 degrees to +/-317 degrees by relocating the azimuth limit switches.

4) A permanent thermal radiation shield was fabricated and installed on the back side of the 2.5m telescope secondary mirror. New counterweights to properly balance the telescope are being fabricated and will be installed in the January bright time.

5) We installed a new switch-over system for the imager UPS power supply. This new system will allow us to switch the imager to a backup UPS should the main UPS unit fail. This task is not quite finished since the last element of the circuit cannot be installed until the imager is powered down. This will occur during the summer 2005 shutdown.

6) During maintenance work, we discovered that the gauge of wire used to power the PT slit controller was undersized for the current load. We re wired this controller with heavier gage wiring and eliminated a potential fire hazard.

7) At the beginning of the quarter we experienced a series of spontaneous reboot problems in the M1 Galil controller. We learned from the controller manufacturer that the problem was most likely caused by poor contacts between the controller and its power supply. We cleaned the contacts and re-installed the problematic controller, and have not had a problem since.

8) We had a series of rotator drive aborts during the quarter. We examined the drive motor for possible slip problems and the rotator for possible spikes in its torque, but found no smoking guns. A closer inspection of the TPM logs over the past few years showed that the frequency of rotator abort problems seems to vary on a yearly cycle, which implies some correlation with temperature. The intermittent nature of the problem makes it difficult to troubleshoot; notwithstanding, we continue to look for the source on a best-effort basis.

9) Seeing studies were made at APO using the data from the 2.5m telescope and the DIMM. The purpose was to determine whether the DIMM measures overall site well enough if it points only at Polaris. Preliminary results indicate that the DIMM tracks seeing reasonably well, with an expected offset. There does appear to be a weak trend with wind speed and direction. Further analysis is required before we can conclusively determine whether leaving the DIMM parked on Polaris is sufficient or whether it is necessary to put effort into making the DIMM capable of bouncing around the sky to obtain a more accurate measurement of site seeing.

3.3. The Photometric Telescope

In the Q3 report, we noted that we had replaced the closed-cycle refrigerator (the CryoTiger) that cools the CCD on the PT with the refurbished spare unit. With the spare in place, we sent the primary unit back to the vendor for refurbishment, given the number of operating hours on the unit. Unfortunately, with the primary unit off-site, we began to see cold head temperature variations in the spare unit installed on the telescope. The magnitude of the variations indicates contamination in the gas circuit of the “refurbished” spare. We plan to re-install the primary unit when it is returned from vendor and send the spare unit back for inspection and repair, since the refurbishment work is still under warranty.

In the Q3 report, we noted the presence of a filament on the surface of the PT CCD. The filament remained in place throughout the quarter, which allowed us to flat-field out the filament from the PT corrected frames. Until the filament becomes problematic, we will leave it in place, as removal requires warming the system and disassembling the camera.

3.4. Operations Software and the Data Acquisition System

All observing software remains under formal version control and all changes are reviewed and approved before work is done. The following work was done on observing software in Q4:

- A new version of IOP was implemented. IOP v3_142_0 included changes to stripeAxisTools to accommodate SEGUE scans, as documented in PR 6067.
- A new version of astrom was implemented and tested; v3_8b fixes the operational problem of goStare writing the opConfig file before goDrift is executed (see PR 6227).
- The MCP and interlock PLC code were updated to reflect the change in the position of the telescope azimuth soft limits. Due to dependencies, this also required updating IOP to v3_143_0 and the watcher/interlock display to v2_30_0.

- . Upgrades were made to the Telescope Performance Monitor (TPM).
  - An array initialization problem in the routine that filters serial link communication status was fixed (see PR4494).
  - A new serial interface to the imager UPS was implemented to record and report alarms generated by the imager UPS system (see PR 6111)
  - The TPM channels that refer to TCC position-velocity-time (PVT) triplets now contain the information from the (struct PVT_M68K) *tccmove* structure in dataCollection.h, rather than duplicating the (struct PVT_M68K) PVT structure that holds the MCP PVTs (see PR 6245).

The production version of the TPM at the end of Q4 was v2_46_0.

4. DATA PROCESSING AND DISTRIBUTION

4.1. Data Processing

4.1.1. Pipeline Development and Testing

No changes were made to the photometric pipeline, Photo, or the spectroscopic pipeline, Spectro1D, in Q4. Work continues on a best-effort basis on the spectroscopic pipeline, idlspec2d, to incorporate the data model changes associated with the public release of DR4.

4.1.2. Data Processing Operations

Data processing operations continue to run very smoothly. In Q4, we processed data from 47 imaging runs, 55 spectroscopic plates, and 165 unique PT patches.

Quality assurance tests were performed on the imaging data and the results posted on the runQA web page. Quality assurance inspections were made on the spectroscopic data and the results summarized on the spectro QA web page. In addition, during observing periods, daily e-mail messages that summarized the results of the most recent reductions were posted by the DP group. These messages provide prompt feedback to the observing and engineering teams regarding the quality of recently acquired data and they alert people to potential problems.

In addition to processing data and preparing for data releases, data processing personnel are involved with the implementation and testing of data reduction software for the SEGUE program; and in the development, installation and commissioning of the Supernova data reduction system at APO. Work performed in these areas is discussed in Section 8.

The DP group continues to back up critical SDSS data into the Enstore tape robot at Fermilab. Critical backups in Q4 included the following:

- Monthly incremental backups of the operational database (opdb);
- All 142 of the DR3 chunks, 71 each for target and best;
- All 168 of the DR4 chunks, 84 each for target and best.

4.2. Data Distribution

4.2.1. Data Usage Statistics

To date, the general public and the astronomy community have access to the EDR, DR1, DR2, and DR3. Figure 4.1 shows the volume of data transferred monthly from the DAS. A total of 930 GB were transferred in Q4, down substantially from previous quarters. Monthly data transfers are characterized by one or two large transfers by institutions apparently transferring much if not all of the data set, and numerous other small transfers. Over the past few months, we have seen a reduction in the number of particularly large transfers.

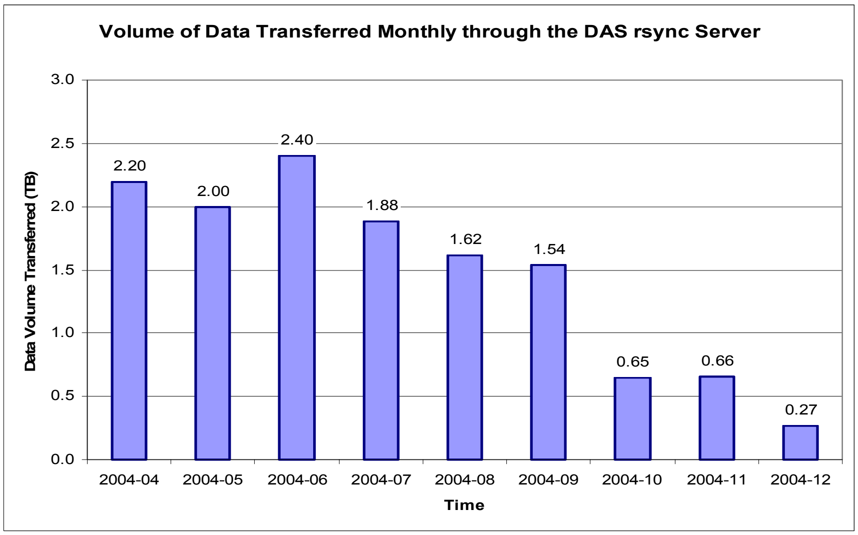

Figure 4.1. Monthly data transfer through the DAS rsync Server

Figure 4.2 plots the number of web hits we receive per month through the various SkyServer interfaces. In Q4 we recorded a total of 8.5 million web hits, compared to 8.2 million web hits in Q3 and 7.4 million hits in Q2. Through December 31, 2004, the various SkyServer interfaces have received a total of 65 million web hits and have processed over 9.3 million SQL queries. As the graph shows, the rate at which the user community is accessing SDSS data continues to grow which

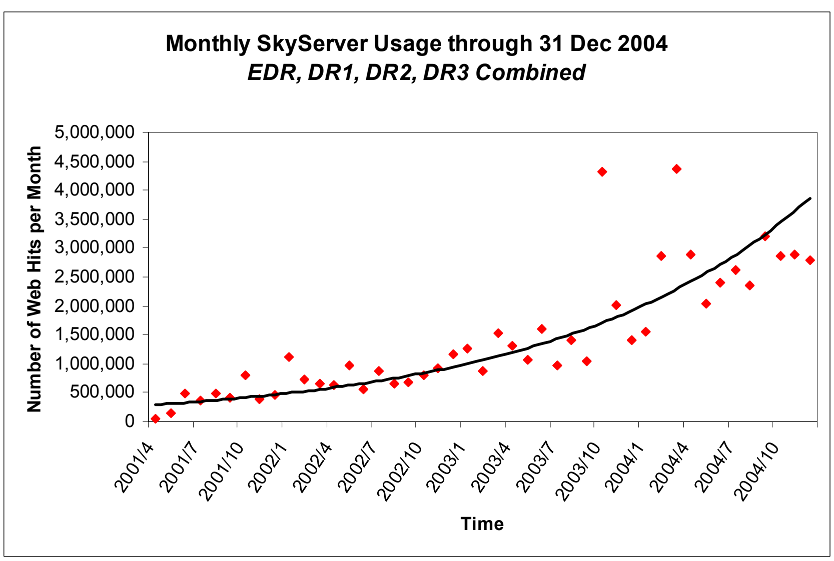

Figure 4.2. SkyServer usage per month, for all public releases combined.

We also measure usage by the number of SQL queries executed. Figure 4.3 shows the total number of SQL queries executed per month. We recorded 1.93 million queries in Q4, compared to 1.84 million in Q3 and 1.73 million in Q2.

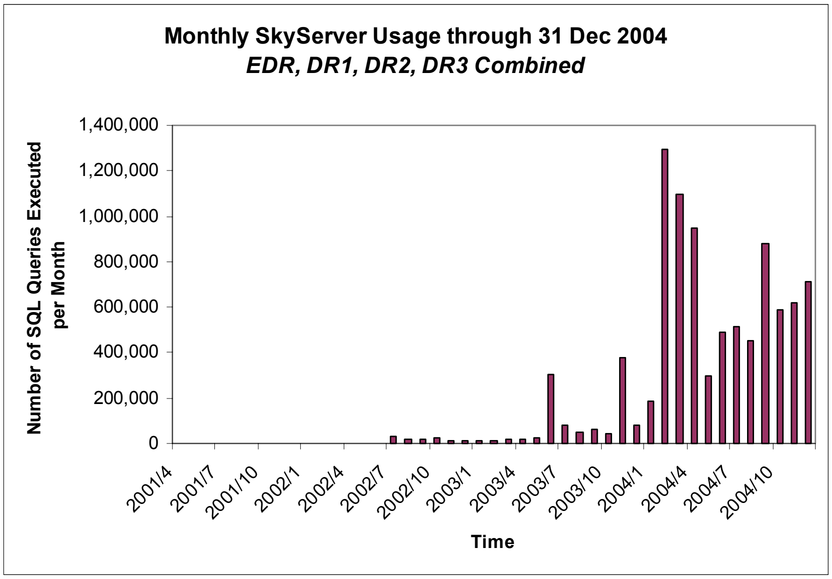

Figure 4.3. SkyServer usage, measured by the number of SQL queries submitted per month.

4.2.2. Data Release 4 Preparations

On December 22, 2004, we made DR4 available to the SDSS collaboration. This provides the collaboration with slightly more than 6 months of proprietary access to the data prior to the public release in July 2005. DR4 is an incremental load on top of DR3 and contains all survey quality imaging data collected through July 2004, and the corresponding spectra. Table 4.1 summarize the contents of DR4 and shows the increment over DR3.

Table 4.1. DR4 Contents

	DR3	*DR4*	*Increment*
Imaging			
Footprint Area	5,282 sq. deg	6,670 sq. deg.	1,388 sq. deg.
Imaging Catalog	141 million objects	180 million objects	39 million
Data volume			
Images	6.0 TB	7.5 TB	1.5 TB
Catalogs (DAS, fits format)	1.2 TB	1.5 TB	0.3 TB
Catalogs (CAS, SQL database)	2.3 TB	3.0 TB	0.7 TB
Spectroscopy			
Spectroscopic Area	4,188 sq. deg.	5,320 sq. deg	1,132 sq. deg.
Total Number of Spectra	528,640	672,640	144,000
Galaxies	374,767	478,887	104,120
Quasars (redshift < 2.3)	45,260	56,340	11,080
Quasars (redshift > 2.3)	5,767	7,478	1,711
Stars	50,369	62,401	12,032
M Stars and later	20,805	26,837	6,032
Sky Spectra	26,819	34,022	7,203
Unknown	4,853	6,675	1,822

In preparation for the DR4 release, web pages describing the contents of DR4 and caveats associated with the data were developed. Existing documentation pages were updated as appropriate, and work was completed on the DAS and CAS, as described in the following sections.

4.2.3. Data Archive Server

The Data Archive Server (DAS) and its associated interfaces continue to be stable and reliable. The DAS is currently serving up the DR1, DR2 and DR3 to the public and collaboration; and DR4 to the collaboration. There were no file corruption problems or unscheduled downtime associated with the DAS in Q4.

Work on the DAS was associated with the DR4 collaboration release. The DAS web pages were updated and new links generated for the new data files associated with DR4. With the web pages and links in place, the various DAS interfaces and data retrieval interfaces were tested and verified ready for use.

4.2.4. Catalog Archive Server

Work on the CAS was associated with finishing the DR4 load for the collaboration. A number of problems were encountered while running the "finish" step of the loading process; these problems greatly increased the time required to complete the loading process. We encountered problems dealing with data that had inadvertently been loaded twice. The process of identifying and removing duplicate objects, and verifying that all references, indices, etc, associated with the data had been properly removed, proved to be tedious and time-consuming. It took a great deal of time and effort to get it right. In addition, we experienced a hard disk failure while running the "finish" step on the TARGDR4 database. Recall that TARGDR4 contains the target version of the data and BESTDR4 contains the best version of each object in our data set. We recovered from the drive failure by replacing the drive, but somewhere in the process some of the data in the partially-finished database became corrupt. We invested a substantial amount trying to repair the data, with only moderate success. In the end, due to schedule constraints, we abandoned the TARGDR4 database load and focused on getting the BESTDR4 database ready for collaboration use.

By the end of December, BESTDR4 was finished, tested, and made available to the collaboration for use and testing. Since over 99% of users access only the "best" version of the data, we felt that the lack of a TARGDR4 database did not warrant holding up the collaboration release of the DAS or BESTDR4.

Given the problems associated with the DR4 load, we plan to reload BESTDR4 and a new version of TARGDR4 in 2005-Q1. We understand why duplicate data was loaded and want to make sure that we can re-do the data load without loading duplicate objects. We are also in the process of installing a new dual-Opteron server that will serve as our new production database loading machine. The new machine uses a RAID10 disk array, which should preclude the type of problems we had with the disk failure during the TARGDR4 load. In short, we want to make sure that we can complete an incremental data load smoothly, from end to end, before making any data model changes and doing a spectro data reload for the public release. We are also in the process of developing and implementing tools and procedures that will verify the integrity of the data at various points in the loading process, and will allow us to more accurately purge and reload corrupt or suspect data. We are also implementing a more stringent version control structure in our production operation to better manage and track code changes. We anticipate that a more structured loading environment, along with the implementation of these tools, will improve the efficiency of future data loads.

In parallel with the database loading effort, development work continued on the code to integrate sector computations and "match" tables into the database. Work was also done on data partioning, which should improve overall performance. At the quarter's end, most of the development work is complete; testing, validation, and final tuning is now underway.

5. SURVEY PLANNING

5.1. Observing Aids

Several programs are used to aid in planning and carrying out observations. Minor bug fixes were made to the plate inventory database.

5.2. Target Selection

For this quarter, 76 plates were designed and drilled in two drilling runs. Of these, 34 were for the Northern survey area and 42 were drilled for the SEGUE program. Of those plates in the Northern survey area, all were located in the area of the "gap" that is the highest priority for observing.

Plate storage at APO continues to be tight. To help alleviate this, plates are occasionally retained at UW after they are drilled until they are actually needed for observing, and two drilling runs will be skipped.

5.3. Survey planning

The software that is used to track survey progress that is contained in this report is also used to prepared monthly observing plans. No significant changes were made this quarter.

6. COST REPORT

The operating budget that the Advisory Council approved for the year 2004 consists of $1,806K of in-kind contributions from Fermilab, US Naval Observatory (USNO), Los Alamos National Laboratory (LANL), and the Japan Participation Group (JPG); and $3,400K for ARC funded expenses.

Table 6.1 shows the actual cost performance by project area for ARC-funded cash expenses in Q4. A more complete table comparing actual to baseline performance is included as an attachment to this report.

Table 6.1. Q4 Cash Expenses and Forecast for 2004 ($K)

	2004 – 4th Quarter		2004 – Total	
Category	Baseline Budget	Actual Expenses	Baseline Budget (Nov 2003)	Current Forecast
1.1. Survey Management	55	51	294	264
1.2. Collaboration Affairs	4	0	16	5
1.3. Survey Operations				
1.3.1. Observing Systems	158	116	648	609
1.3.2. Data Processing & Dist.	137	103	593	603
1.3.3. Survey Coordination	0	0	0	0
1.3.4. Observatory Support	380	405	1,522	1,559
1.4. ARC Corporate Support	33	15	176	63
Sub-total	767	690	3,248	3,104
1.5. Management Reserve	32	0	152	0
1.6. Capital Improvements	0	0	0	43
Total	799	690	3,400	3,147

6.1. Q4 Performance - In-kind Contributions

The sum of in-kind contributions for the fourth quarter was $299K against the baseline forecast of $450K and was provided by Fermilab, Los Alamos, and the U.S. Naval Observatory (USNO).

Fermilab provided telescope engineering and maintenance support, and the data processing systems at Fermilab, as agreed. Details of Q4 in-kind contributions are as follows:

- As reported in previous reports, the level of in-kind engineering support at APO was less than the baseline forecast because the engineer who serves as the SDSS Telescope Engineer became involved in another Fermilab astrophysics project in Q1. As a result, his level of support for the SDSS project has been reduced from 100% to 50%.
- The level of in-kind effort for Software and Data Processing support at Fermilab is less than the baseline forecast; the actual level of effort required to support on-going operations is less than we had anticipated would be required when the budget was prepared a year ago.

Los Alamos provided programming support for the Telescope Performance Monitor and assisted with the preparation of a technical paper documenting the Photometric Telescope data reduction pipeline, MTPIPE. The level of in-kind infrastructure support provided in Q4 was substantially less than the baseline forecast but sufficient to complete the necessary work.

USNO provided support as required for the astrometric pipeline and other software systems they maintain. As previously reported, the value of in-kind support was lower than the baseline because the baseline did not properly reflect the lower level of anticipated support in 2004.

6.2. Q4 Performance – ARC Funded Expenses

ARC-funded expenses in Q4 were $690K, or $77K (10%) below the fourth quarter budget of $767K, excluding management reserve.

Survey management costs were $52K against a Q4 budget of $55K (-7%). Actual Q4 expenses for the various management-related SSP accounts were within a few $K of their respective budgets. The plus/minus variations cancelled each other out for the category as a whole. For the year, the forecast for Survey Management is $264K, or $30K (10%) below the baseline budget. The largest variations were as follows:

- Support for Survey Management at Fermilab was overspent because computing hardware to support the data distribution effort was charged to this account. A new account has been established to properly capture these costs in 2005.
- The budget for Public Affairs was under spent. We had budgeted $10K for data release brochures but decided during the year that such brochures were not necessary. AAS meeting expense costs were also less than anticipated.
- Actual expenses for the Public Information Officer were significantly less than anticipated.

The budget for Collaboration Affairs provides for Working Group travel and technical page charges and is held in an ARC corporate account. No expenses were incurred in Q4. For the year, the forecast for Collaboration Affairs is $5K, or $11K (69%) below the baseline budget of $16K.

Observing Systems costs were $116K against a Q4 budget of $158K (-26%). Actual Fermilab and Princeton costs for Observing Systems Support were within a few $K of their respective budgets. Actual UW costs were significantly less than budgeted; the amount of off-site engineering and technical effort required to support on-going operations was less than we had anticipated when the budget was prepared. Funds from the ARC Observing Systems Support account were used to pay the balance due on primary mirror aluminizing, settle final relocation costs for one of our observers, and cover miscellaneous travel and teleconference expenses associated with fall observing operations. For the year, the forecast for Observing Systems is $610K, or $38K (6%) below the baseline budget of $648K.

Data Processing and Distribution costs were $103K against a Q4 budget of 137K (-24%). Fermilab expenses exceeded the budget slightly due to computer hardware purchases associated with data distribution operations. Princeton expenses were substantially less than the baseline due to a credit for effort that was incorrectly charged to the account earlier in the year. JHU costs were slightly less than budgeted and UC costs were in close agreement with the budget. For the year, the forecast for Data Processing and Distribution is $603K, or $10K (2%) above the baseline budget of $593K.

Observatory Support costs were $405K against a Q4 budget of $380K (6%). Salary costs exceeded the baseline slightly because we added the part-time observer to our staff after the baseline budget was prepared. Other observatory operating expenses were in line with expectations. For the year, the forecast for Observatory Support is $1,559K, or $37K (2%) above the baseline budget of $1,522K.

Miscellaneous ARC corporate expenses (i.e., audit fees, bank fees, petty cash, and APO trailer rentals) were as expected. No charges were incurred against the observers' development fund or the Additional Scientific Support fund in Q4. For the year, the forecast for ARC Corporate Support is $63K, or $113K (64%) below the baseline budget of $176K.

6.3. Management Reserve

No management reserve funds were expended in 2004. Given estimated final cash expenses of $3,147K, we anticipate moving $253K in unspent funds into 2005 to complete the survey and pay down unpaid invoices. The final amount moved forward will be established once final 2004 invoices are received and paid.

7. PUBLICATIONS

Through December 31, 2004 there are 633 published refereed papers that include 'SDSS' or 'Sloan' in their title and/or abstract. These papers have been cited a total of 13,143 times, including 23 papers cited more than 100 times and 64 papers with 50 or more citations. In 2004, over half the SDSS papers were written by people outside the SDSS collaboration, meaning that the broader astronomical community is making good use of the public data releases.

The York et al (2000) paper has been cited over 500 times, which puts it in the category of "renowned" papers. The Tegmark et al. (2004) paper on cosmological parameter estimation is the single most cited paper in all of astrophysics, published in 2004 (SPIRES gives it over 300 citations). The number of SDSS papers is growing exponentially, with a doubling time of 1.5 years; the number of citations to those papers is also growing exponentially, with a doubling time of 1.2 years.

In addition, there have been 921 un-refereed papers with "SDSS" or "Sloan" in the title and/or abstract.

The following papers were submitted and/or published in Q4:

SDSS galaxy-galaxy weak lensing: testing redshift distributions and other systematics
MNRAS submitted - Rachel Mandelbaum, et al

The Sloan Digital Sky Survey Quasar Catalog III. Third Data Release
AJ submitted - Donald P. Schneider, et al

Baryonic signature in the large-scale clustering of Sloan Digital Sky Survey quasars
PASJ submitted - Kazuhiro Yahata, et al

Photometric Accretion Signatures Near the Substellar Boundary
AJ submitted -Peregrine M. McGehee, et al

Optically Identified BL Lacertae Objects from the Sloan Digital Sky Survey
AJ submitted - Matt Collinge, et al

The SDSS u-band Galaxy Survey: Luminosity functions and evolution
MNRAS accepted -Ivan Baldry, et al

Cosmic homogeneity demonstrated with luminous red galaxies
ApJ submitted - David W. Hogg, et al

Rotation Velocities of Two Low Luminosity Field Galaxies
ApJL submitted - Jim Pizagno, et al

A comprehensive model for the Monoceros tidal stream
ApJ submitted - Jorge Penarrubia, et al

Cross Correlating the CMB with LSS : The ISW Effect
PRD submitted - Nikhil Padmanabhan, et al

A New Milky Way Companion: Unusual Globular Cluster or Extreme Dwarf Satellite?
AJ submitted - Beth Willman, et al

SDSS~J210014.12+004446.0: A new Dwarf Nova with Quiescent Superhumps?
PASP accepted - Jonica Tramposch, et al

The Small-scale Clustering of Luminous Red Galaxies via Cross-Correlation Techniques
ApJ accepted - Daniel Eisenstein, et al

Publications Based on Public Data

Sloan Digital Sky Survey Quasars in the SWIRE ELAIS N1 Field: Properties and Spectral Energy Distributions
AJ submitted – Evanthia Hatziminaoglou, et al

Galaxy Groups in the SDSS-DR3
ApJ submitted - Manuel Merchan Ariel Zandivarez

The scale-dependence of relative galaxy bias: encouragement for the "halo model" description
ApJ submitted - Michael R. Blanton, et al

Photometric Determination of R$_{200}$ for SDSS Galaxy Clusters
ApJ submitted - Sarah Hansen, et al

Cool White Dwarfs in the Sloan Digital Sky Survey
AJ submitted - Mukremin Kilic, et al

On The Mass-to-Light Ratio of Large Scale Structure
ApJ submitted - Jeremy Tinker, et al

On the Incidence and Kinematics of Strong Mg II Absorbers
ApJ submitted - Gabriel E. Prochter, et al

Systematics of the Ultraviolet Rising flux in a GALEX/SDSS sample of Early-type Galaxies
ApJL accepted - R. M. Rich, et al

Structure Function Analysis of Long Term Quasar Variability
AJ accepted - W. H. de Vries, et al

The Ultraviolet Luminosity Function of GALEX Galaxies at Photometric Redshifts Between 0.07 and 0.25
ApJL accepted - T. Budavari, et al

Toward Understanding Environmental Effects in SDSS Clusters
A&A submitted - J. Einasto, et al

Evolution of the Cluster Mass and Correlation Functions in LCDM Cosmology
ApJ accepted - Joshua D. Younger, et al

Chandra-SDSS Normal and Star-Forming Galaxies I: X-ray Source Properties of Galaxies Detected by Chandra in SDSS DR2
AJ accepted - A.E. Hornschemeier, et al

8.0. SDSS-II DEVELOPMENT WORK

8.1. Progress on the SEGUE Project

SEGUE observing (both imaging and spectroscopy, depending on photometricity) was interleaved with SN repeated imaging during the period October through mid-November 2004. After November 17, 2004, SEGUE observing was obtained when Legacy NGC plates or imaging was not available due to hour angle or airmass constraints.

A total of about 300 non-overlapping square degrees of usable SEGUE quality imaging was obtained during the quarter, and 19 SEGUE plates (corresponding to about 26 plate equivalents as seguefaint plates are of longer exposure) were completed.

8.1.1. SEGUE Imaging

The SEGUE imaging obtained in Q4 included some at relatively southern declinations (that along stripe 90), near dec = -20 deg, near the limit of what the SDSS system can do without significant degradation of image quality.

Short SEGUE runs (of duration < 1 hour), frequently didn't cross an existing PT patch, and so the calibration for them will be deferred until PT calibration patches can be obtained. One of 11 scans was declared non-photometric and will be rescanned if time permits in future years.

8.1.2. SEGUE Spectroscopy

The SEGUE spectroscopy obtained in Q4 was the first to test the SEGUE target selection v3_0 code and the first to test the longer faint plate exposure S/N^2 limits, which were empirically determined last quarter to be somewhat too short. The new faint plate exposure guidelines are for the observers to integrate until the $(S/N)^2$ recorded on the mountain for the faint plate for all cameras exceeds 50 (previous limit 30). And preliminary inspection and reductions of these faint SEGUE plates indicates that this increased exposure is sufficient to obtain metallicity parameters for faint stars not available at lesser exposures. We will continue to operate with these limits for the faint plates, pushing our effective plate-equivalents per plate-pair ratio to 3.0, up from 2.6.

Additionally, SEGUE obtained exposures this quarter of three globular clusters of known metallicity and gravities, and successfully demonstrated that the stellar atmosphere parameter measuring code techniques of Beers/Allende-Prieto/Wilhelm/Norris can recover the metalicities of these stars to accuracy of better than one sigma = 0.3 dex for stars of relatively high S/N, a main technical goal of the SEGUE science program.

8.1.3. SEGUE Progress To Date:

To date, SEGUE has obtained approximately 1050/3900 square degrees of imaging and approximately 56/400 plates (78/600 plate equivalents) towards its baseline goal.

8.2. Progress on the Supernova Project

The early science and test/engineering run of the SDSS II Supernova Survey took place this quarter. From late September through mid-November, 20 nights of SDSS 2.5m time were scheduled for imaging on stripe 82N. Due to weather, about half of these turned out to be scientifically useful.

The imaging data were processed through the compute cluster set up at APO for this purpose, using 7 dual processor machines and a RAID array for storage. Data were spooled from tape, processed through PHOTO up to corrected frames, processed through a frame subtraction pipeline (g and r only), and placed into a database, where object positions were matched in g and r (to veto asteroids) and matches with known variables (stars and AGN) were also vetoed. The resulting candidates were then inspected by eye using ds9 and ranked by the SN team for possible spectroscopic follow-up. Promising candidates were displayed on a web page and were followed up spectroscopically using the ARC 3.5m, the HET, and, in one case, the SN Factory using the UH 88-in. Additional imaging follow-up using the ARC 3.5m and the NMSU 1m, especially after the SDSS run ended, proved quite useful. Generally, the processing went smoothly, though there were occasional hardware glitches; in addition, since this was a test run, parts of the pipeline were updated during the run, though the bulk of development work upstream of the handscanning environment was done in Q3. When the system was working well, a full night of 6-column data could be processed through the system in 50-60 hours, close to our goal of 48 hours. Since there are 3 stripe 82N template image regions and these were processed sequentially, we were often starting human inspection of candidates within ~30-40 hours.

During the 2.5m run, at least one and often two or three members of the SN team were at the mountain, processing the data or working on the 3.5m follow-up. The team met daily by phonecon. The APO observing and computer support staff was also critical to the success of the run.

The test run produced 16 spectroscopically confirmed supernovae of type Ia, 5 confirmed type II, and 1 type Ibc. Several other type Ia's are probable but were not spectroscopically confirmed. A substantial number of these supernovae have publication-quality lightcurves (and will be published). Near the end of Q4, a pipeline for precision photometry of these supernovae was under development. In addition, the team was studying the usefulness and quality of the data that was either non-photometric and/or moony; preliminary indications are that data extending significantly into bright time will be scientifically useful and, according to our Monte Carlo simulations, necessary to achieving the desired number of high-quality SN lightcurves.

In Q4, post-run studies also commenced of our SN detection efficiency. While the number of quality lightcurves we obtained in the test run was about that predicted by the simulations, the total number of SN detected was lower than expected particularly beyond $z \sim 0.25$. Subsequent tests of the frame subtraction code, examination of the i-band test run data, and tests on supernovae found in Fall 2002 SDSS data, have revealed that the parameters set in the object identification code in the frame subtraction pipeline were improperly weeding out faint objects in the subtracted frames. In SDSS II, we will correct this, add real-time i-band subtraction, and also introduce artificial supernovae into the data stream to monitor detection efficiency in real time.

Exhibit 1. CY2004 Actual Cash Expenses (Preliminary)

SDSS CY2004 Budget Forecast as of December 31, 2004 (*in $000s*)

		Qtr 1 Jan-Mar	Qtr 2 Apr-Jun	Qtr 3 Jul-Sep	Qtr 4 Oct-Dec			CY2004 Budget Forecast		
		Actual Expenses	Actual Expenses	Actual Expenses	Approved Budget	Actual Expenses	Variance (%)	Approved Budget	Final Expenses	Variance (%)
ARC-FUNDED BUDGET										
1.1 Survey Management										
SSP21	ARC Secretary/Treasurer	2	2	2	3	2	-31%	12	9	-22%
SSP34	ARC Business Manager	10	19	14	15	15	-1%	58	58	0%
SSP46	PU Office of the Project Scientist	1	.2	55	8	2	-69%	59	58	-1%
SSP48	FNAL Support for Survey Management	12	14	18	9	10	7%	36	54	49%
SSP67	UC Support for Survey Management	12	12	30	14	15	5%	79	69	-12%
SSP91a	ARC Support for Public Affairs	2	.2	0	0	3	----	24	5	-80%
SSP91i	ARC Support for Public Information Officer	4	2	2	6	4	-32%	27	12	-57%
	Sub-total	42	49	121	55	51	-7%	294	264	-10%
1.2 Collaboration Affairs										
SSP91c	ARC Support for Collaboration Affairs	2	1	2	4	0	-100%	16	5	-69%
	Sub-total	2	1	2	4	0	-100%	16	5	-69%
1.3.1 Observing Systems										
SSP42	FNAL Observing Systems Support	62	48	75	38	36	-4%	149	221	48%
SSP61	FNAL Observing Programs and DA Support	1	3	2	5	3	-46%	23	9	-64%
SSP31	UW Observing Systems Support	66	57	88	72	45	-37%	288	257	-11%
SSP32	PU Observing Systems Support	10	12	18	15	18	23%	59	59	-1%
SSP91d	ARC Observing Systems Support	6	24	21	29	14	-52%	128	64	-50%
	Sub-total	145	145	204	158	116	-26%	648	610	-6%
1.3.2 Data Processing and Distribution										
SSP40	FNAL Software and Data Processing Support	48	54	49	27	33	23%	128	185	44%
SSP38	PU Software and Data Processing Support	56	47	24	33	1	-97%	148	128	-13%
SSP39	UC Software and Data Processing Support	9	10	10	11	10	-2%	42	40	-6%
SSP37	JHU Data Archive Development and Support	65	51	51	57	49	-14%	240	215	-10%
SSP33	UC Operations Support	8	9	9	9	9	8%	35	35	2%
	Sub-total	186	171	143	137	103	-24%	593	603	2%
1.3.4 Observatory Support										
SSP35	NMSU Site Support	384	391	380	380	405	6%	1,522	1,559	2%
	Sub-total	384	391	380	380	405	6%	1,522	1,559	2%
1.4 ARC Corporate Support										
SSP91e	ARC Corporate Support	10	17	14	12	15	23%	62	55	-11%
SSP91f	ARC Additional Scientific Support	0	0	0	18	0	-100%	102	0	-100%
SSP91h	ARC Observers' Research Support	0	4	4	3	0	-100%	12	8	-36%
	Sub-total	10	21	18	33	15	-56%	176	63	-64%
SUBTOTAL		769	778	868	767	690	-10%	3,248	3,104	-4%
SSP91g	Capital Improvements	0	43	0	0	0	---	0	43	---
SSP91	Management Reserve[1]	0	0	0	32	0	-100%	152	0	-100%
TOTAL ARC-FUNDED BUDGET		769	778	868	799	690	-14%	3,400	3,147	-7%

Exhibit 2. CY2004 In-Kind Contributions (Preliminary)

SDSS CY2004 Budget Forecast as of December 31, 2004 (*in $000s*)

		Qtr 1 Jan-Mar	Qtr 2 Apr-Jun	Qtr 3 Jul-Sep	Qtr 4 Oct-Dec			CY2004 Budget Forecast		
		Actual Expenses	Actual Expenses	Actual Expenses	Approved Budget	Actual Expenses	Variance (%)	Approved Budget	Final Expenses	Variance (%)
IN -KIND CONTRIBUTION										
1.1 Survey Management										
SSP48	FNAL Support for Survey Management	42	41	38	49	38	-23%	191	159	-17%
	Sub-total	42	41	38	49	38	-23%	191	159	-17%
1.3.1 Observing Systems										
SSP42	FNAL Observing Systems Support	50	52	55	68	42	-39%	267	198	-26%
SSP58	LANL Observing Systems Support	59	40	30	57	39	-31%	222	167	-25%
SSP61	FNAL Observers' Programs and DA Support	5	2	7	13	0	-100%	52	14	-74%
	JPG Observing Systems Support	0	0	0	0	0	---	0	0	---
	Sub-total	114	93	91	138	81	-42%	541	378	-30%
1.3.2 Data Processing and Distribution										
SSP40	FNAL Software and Data Processing Support	293	167	126	229	165	-28%	941	752	-20%
SSP57	USNO Software and Data Processing Support	16	14	17	34	15	-54%	133	63	-53%
	Sub-total	310	181	143	263	181	-31%	1,074	815	-24%
TOTAL IN-KIND CONTRIBUTION		465	316	273	450	299	-34%	1,806	1,352	-25%
TOTAL BUDGET		1,234	1,093	1,141	1,249	989	-21%	5,206	4,500	-14%