Sloan Digital Sky Survey II

2005 FOURTH QUARTER REPORT

October 1, 2005 – December 31, 2005

Table of Contents

Q4 PERFORMANCE HIGHLIGHTS

- We obtained 67 square degrees of new imaging data for the Legacy Survey, against a baseline goal of 125 square degrees. We completed 28 spectroscopic plates against a goal of 19 plates.
- We obtained 313 square degrees of new imaging data for the SEGUE Survey, against a baseline goal of 306 square degrees. We also completed 17 SEGUE plates (8 bright and 9 faint). The combination is roughly equivalent to completing 9 SEGUE tiles, against a goal of 17 tiles.
- We completed the first full observing season for the Supernova (SN) Survey. We obtained 6,675 square degrees of new imaging data against a baseline goal of 5700 square degrees. The data resulted in the discovery of 126 spectroscopically confirmed supernovae of Type Ia, 13 spectroscopically probable Type Ia's, and a smaller number of other supernova types. Our baseline goal was to discover 50 to 60 Type Ia supernovae per season sufficiently early in their development that they are useful for cosmology.
- An initial version of Data Release 5 was made available to the collaboration on October 21, 2005. We intend to release a more complete version, with a small number of improvements, to the collaboration in March 2006. The DR5 public release is scheduled for June 30, 2006.
- We completed the recruitment of three new observers, which brings us back to the planned level of observer staffing for SDSS-II operations.
- Q4 cash operating expenses were $914K against a baseline budget of $869K, excluding management reserve. In-kind contributions were $217K against anticipated contributions of $122K. No management reserve funds were expended in Q4.
- For the year ended December 31, 2005, cash operating expenses were $1,979K against a baseline budget of $2,305K, excluding management reserve. In-kind contributions were $459K against a baseline budget of $355K.

1. SOME RECENT SCIENCE RESULTS

Beginning with this Quarterly Report, we include brief descriptions, with graphics, of some of the scientific work accomplished during the reporting interval (bearing in mind that the effort will often spill over into other quarters). Unlike the list of publications given in Exhibits 3 and 4, the selection of topics here is by no means comprehensive, nor even representative, of the science being undertaken by the SDSS collaboration. These short science descriptions nevertheless augment our reporting of activities in SDSS-II.

A Galaxy So Close We Couldn't See It Before

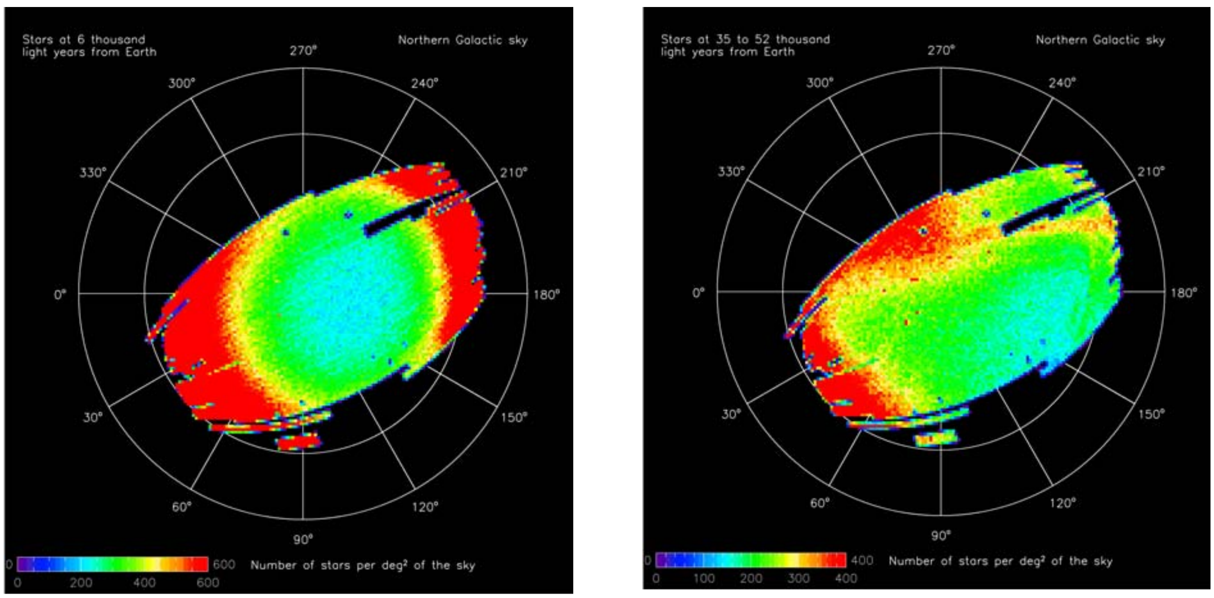

3-dimensional maps of the distribution of distant stars have revealed a remarkable new structure in the northern sky, which appears to be a new dwarf galaxy companion to the Milky Way. The two panels show the surface density of stars on the sky, in spherical shells of distance, over the region covered by the SDSS imaging survey. Distances to the stars can be estimated from the combination of their colors and brightnesses. The left panel shows relatively nearby stars, within 6000 light years. The distribution is symmetric about a horizontal line through the middle of the panel, as expected for a symmetric Milky Way. The right panel shows stars at distances of 35,000 to 52,000 light years. This distribution is remarkably asymmetric, with an extended over-density of stars near the top edge. The most likely interpretation is that this over-density is a dwarf galaxy that is merging with the Milky Way and being disrupted by the Milky Way's gravity. It covers an enormous area of sky, roughly 5000 times the area of the full moon, because it is so close - in fact, it is the largest feature on the sky other than the Milky Way itself. However, the galaxy is one of the intrinsically faintest known, comparable in luminosity to the largest globular star clusters, but much more spread out in space. The diffuseness of this structure made it impossible to pick out from earlier star maps. The precise multi-color photometry of the SDSS allows us to examine the stellar distribution with “3-d glasses," and our new neighbor then snaps into view.

References

1. M. Juric et al., Milky Way Tomography with the SDSS, preprint astro-ph/0510520, submitted to The Astrophysical Journal

Einstein's Bullseyes

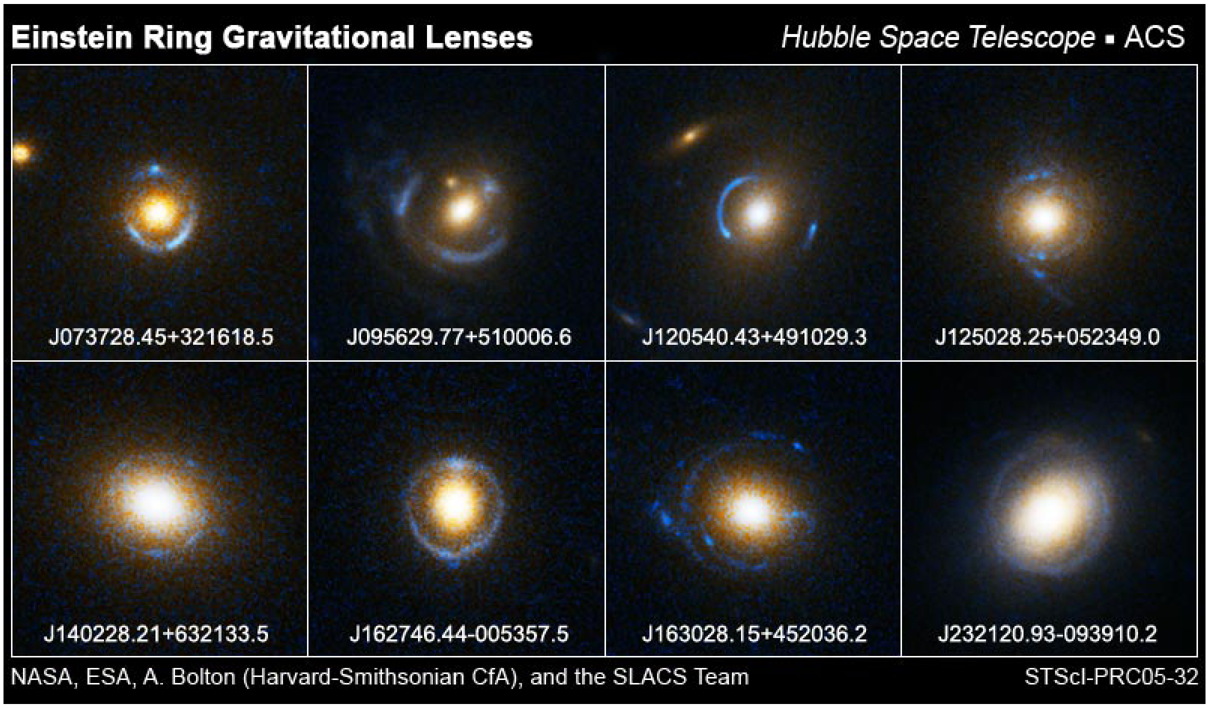

Each of these spectacular Hubble Space Telescope images shows a distant blue galaxy whose light has been distorted into a ring by the gravity of the massive, red foreground galaxy in the center. This gravitational lensing phenomenon was predicted by Albert Einstein in 1935. While observations since the 1970s have uncovered many cases of galaxies producing multiple images of background quasars, the "bulls-eye" phenomenon seen here requires near perfect alignment of the foreground and background galaxy, so it is very rare.

Adam Bolton, Scott Burles, and their collaborators have combed the SDSS to discover these rare jewels by searching the galaxy spectra to find cases where the light of an old, massive galaxy is mixed with the light of a young galaxy at a higher redshift. They then obtained high-resolution, HST images of these systems to search for gravitational lenses (at the ground-based resolution of SDSS images, the foreground and background galaxies would be blurred together). This approach has proved stunningly efficient, revealing 19 strong gravitational lens candidates from 28 observed systems.

Among these 19 are the eight "optical Einstein rings" shown here. Only three such visible light rings were known previously (several have also been found in radio images), so the SDSS/HST survey has nearly quadrupled the number of objects in this class. These lensing systems are useful as well as beautiful, measuring the distribution of invisible dark matter in the lensing galaxies. Recent analysis of these systems shows that the red, elliptical galaxies have extended dark matter halos similar in form to those that surround spiral galaxies like the Milky Way.

References:

1. Bolton et al., The Sloan Lens ACS Survey. I. A Large Spectroscopically Selected Sample of Massive Early-Type Lens Galaxies, preprint astro-ph/0511453, The Astrophysical Journal, in press.
2. T. Treu et al., The Sloan-Lens ACS Survey II: stellar populations and internal structure of early type lens galaxies, preprint astro-ph/0512044, The Astrophysical Journal, in press
3. L. Koopmans et al., The Sloan Lens ACS Survey. III - The Structure and Formation of Early-type Galaxies and their Evolution since z ~ 1, preprint astro-ph/0601628, submitted to The Astrophysical Journal.

2. SURVEY PROGESS

As we have noted in past reports, it is more practical and convenient to report progress and performance metrics in terms of complete observing runs as opposed to partial runs. The period of accounting for this report includes three complete observing runs spanning the period from September 15 through December 15. The time from December 16 through 31 will be accounted for in the next report.

2.1. Legacy Survey

Observing operations associated with the Legacy Survey occurred during the November and December observing runs. Table 2.1 compares the imaging and spectroscopic data obtained against the baseline plan. We obtained 67 square degrees of new imaging data against a baseline goal of 125 square degrees. We completed 28 Legacy spectroscopic plates against the baseline goal of 19 plates. Weather was more suitable for spectroscopy than imaging when the Northern Galactic Cap was visible from APO.

Table 2.1. Legacy Survey Progress in 2005-Q4

	2005-Q4		Cumulative through Q4	
	Baseline	Actual	Baseline	Actual
Legacy Imaging (sq. deg)	125	67	7808	7501
Legacy Spectroscopy (tiles)	19	28	1158	1161

The following graphs show progress against the baseline plan. For the Legacy Survey, we have chosen to extend the progress charts from SDSS to include the three-year time extension. For the imaging survey, the baseline has been left unchanged. For the spectroscopic survey, we have set the baseline plan for SDSS-II equal to actual progress prior to July 2005. In addition to showing the rate at which we need to complete plates to finish the Survey, this shows the rate at which we completed plates in the past. As a point of reference, we have also included the original baseline plan for SDSS-I on the chart.

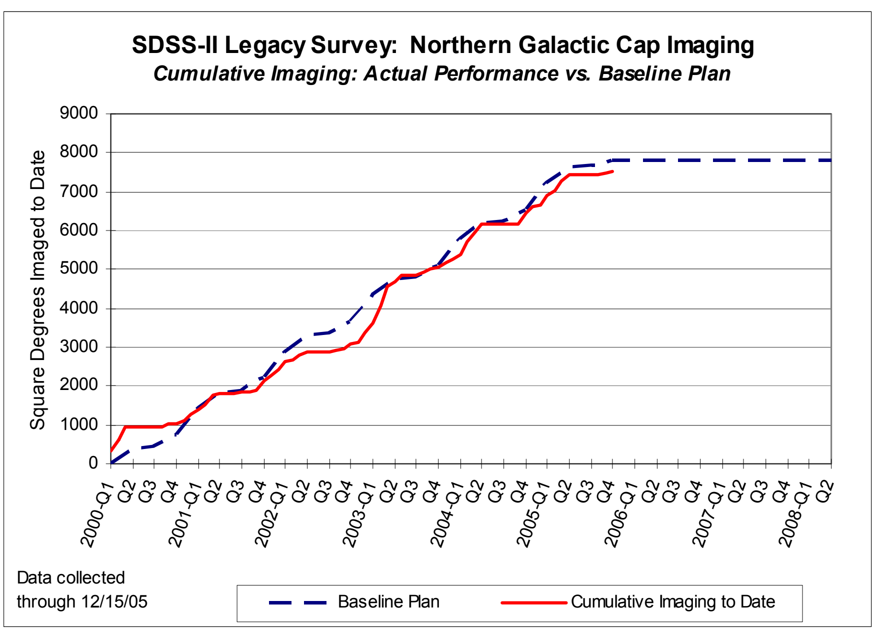

Figure 2.1. Imaging Progress against the Baseline Plan – Legacy Survey

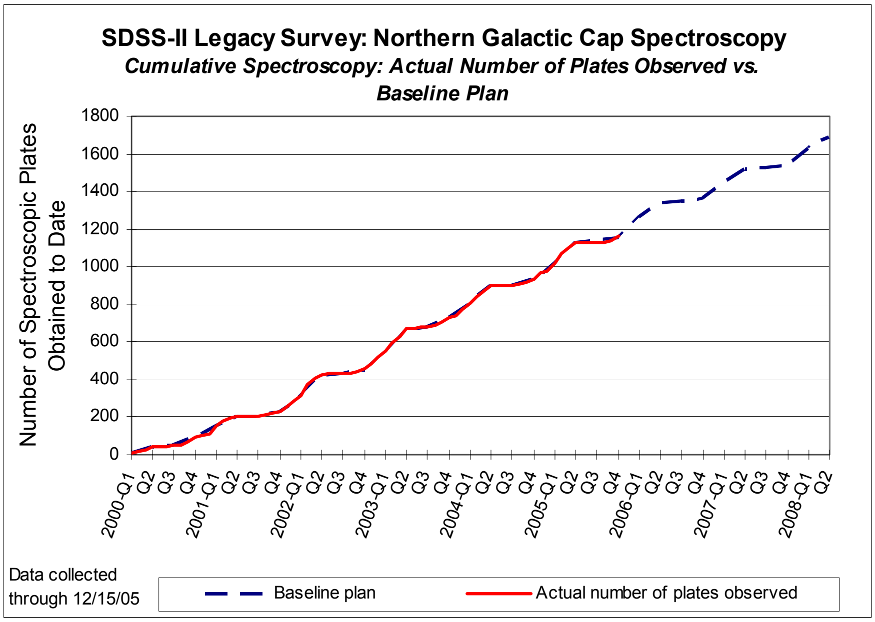

Figure 2.2. Spectroscopic Progress against the Baseline Plan – Legacy Survey

2.2. SEGUE Survey

Table 2.2 compares SEGUE progress against the baseline plan. The SEGUE Survey is ahead of the baseline in both imaging and spectroscopy due to the acquisition of SEGUE data in previous quarters, when commissioning and proof-of-concept observations were made.

SEGUE imaging data was obtained under photometric conditions during the October, November and December observing runs. Observing time in October and November was shared with the Supernova Survey following the time-sharing arrangement outlined in the "Schedule for the Three-Year Baseline Plan." We obtained a total of 313 square degrees of new SEGUE imaging data from stripes 79, 1260, and 1300. Data from stripe 1300 was turned around quickly in order to drill SEGUE plates to be observed in 2006-Q1.

Table 2.2. SEGUE Survey Progress in 2005-Q4

	2005-Q4		Cumulative through Q4	
	Baseline	Actual	Baseline	Actual
SEGUE Imaging (sq. deg)	306	313	507	1649
SEGUE Spectroscopy (bright plates)	17	8	28	55
SEGUE Spectroscopy (faint plates)	17	9	28	44

SEGUE spectroscopic observing occurred during November when we were not imaging for the SN Survey, and during December when we were not imaging for the Legacy and SEGUE Surveys. A total of 17 SEGUE plates (8 bright and 9 faint) were completed. This is roughly equivalent to completing 9 SEGUE tiles, as a SEGUE tile is considered complete when the faint and bright plate combination for a field is observed. Our observing strategy is arranged to complete plate pairs in roughly the same time frame, in order to maximize the scientific usefulness of each plate pair. However, given the many factors that affect observing operations (atmospheric conditions, available time, etc.), it is not always efficient to complete plates in "pair combinations." Therefore, we have elected to separately report progress in terms of the number of bright and faint plates completed, as opposed to combined bright/faint plate pairs (i.e., SEGUE tiles).

The following graphs illustrate SEGUE progress against the baseline plan. The imaging graph presents a straightforward comparison of imaging progress against plan. The graph illustrating spectroscopic progress shows the rate at which we've completed bright and faint plates separately.

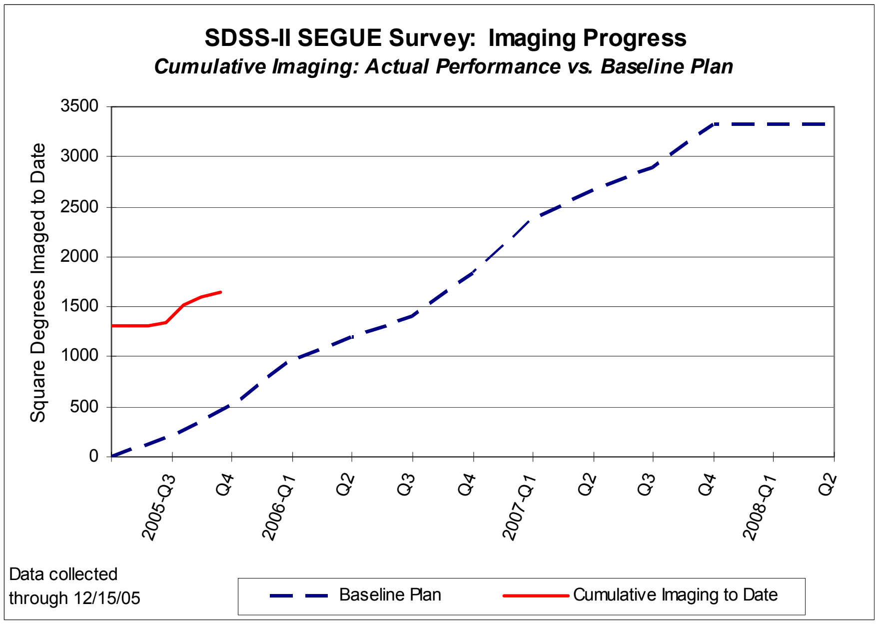

Figure 2.3. Imaging Progress against the Baseline Plan – SEGUE Survey

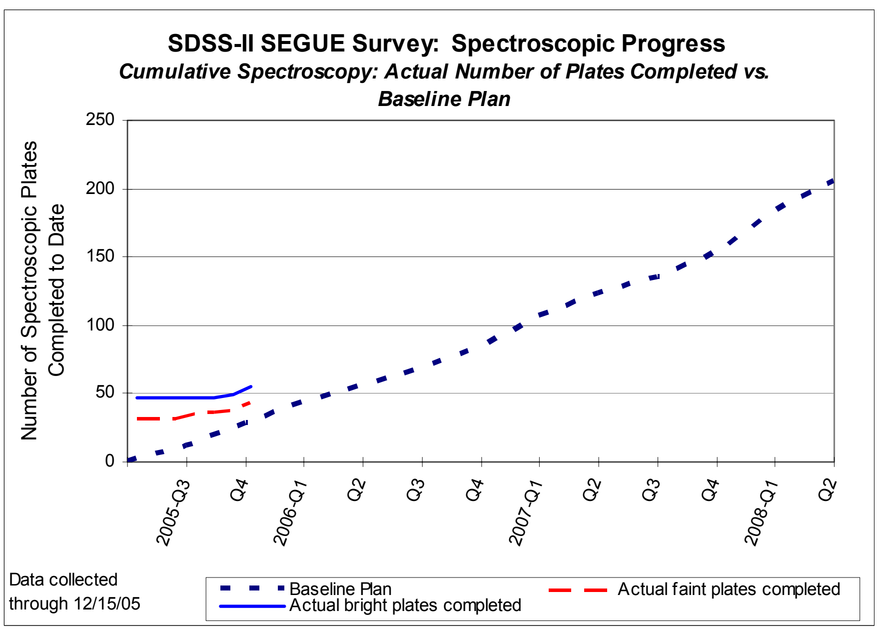

Figure 2.4. Spectroscopic Progress against the Baseline Plan – SEGUE Survey

Finally, in addition to observing at APO, the SEGUE team observed with the USNO 1-m telescope in Flagstaff during dark time to obtain images of bright stars and cluster with the u'g'r'i'z' system, which will be used to refine the calibration of the SEGUE photometry and spectroscopy.

2.3. Supernova Survey

The 2005 observing campaign for the Supernova Survey began on September 1 and ended on November 30. During this period, we obtained 6,675 square degrees of new imaging data on Stripe 82 against a baseline goal of 5700 square degrees. Analysis of the data and follow-up observations resulted in the discovery of 126 spectroscopically confirmed supernovae of Type Ia, 13 spectroscopically probable Type Ia's, and a smaller number of other supernova types. Based on Monte Carlo simulations, our expectations going into the survey were that we would measure ~50-60 type Ia light curves per season with high enough sampling rate and signal to noise to be useful for cosmological studies. While the fall 2005 sample is still being analyzed, it is our expectation that a high fraction of the 126 spectroscopically confirmed Ia’s will in fact be useful for cosmological analysis, exceeding our initial expectations. One indication of this is that a very preliminary Hubble diagram, based on initial reductions of the 126 Type Ia's, shows a dispersion in peak magnitude of about 0.15 mag after corrections for light-curve shape and extinction, comparable to that seen in other recent SN surveys. We also note that host-galaxy spectra (and redshifts) were obtained in December for another ~30 candidates which have light curves strongly indicative of Ia's but for which confirming supernova spectra were not obtained during the run; we expect these objects will be quite useful for science analysis as well.

Nearly all of the observed objects have been released promptly in CBET[1] circulars, and we have a public webpage where they are listed: http://sdssdp47.fnal.gov/sdsssn/snlist.php . The Type Ia redshift range is from 0.01 to 0.42 and the average Type Ia redshift is 0.21, very close to what was expected.

In addition to supporting observing operations, members of the SN team devoted a significant amount of time to commissioning the data processing and analysis systems for the Survey. Team members were responsible for on-mountain data processing at APO and kept the system up to speed despite compute cluster crashes, occasional glitches in the new data acquisition system, and occasional difficulties processing data of substandard quality through the pipelines. These challenges were all overcome and the system was able to keep up with the data rate extremely well. Using artificial supernovae, the team instituted a program to monitor the performance of the on-mountain data system, checking detection efficiency and photometric accuracy as functions of magnitude for each run. Team members also participated heavily in human scanning of thousands of supernova candidates. Upgrades are being planned for the on-mountain software system for the fall 2006 season, based on lessons learned from this first season.

At quarter’s end, the SN team was in the process of preparing a web site for the public distribution of the Stripe 82 data from the 2005 campaign. It is anticipated that the public release will occur in mid-January 2006.

3.0 OBSERVING EFFICIENCY

Observing efficiency is summarized according to the categories used to prepare the baseline projection.

[1] Central Bureau for Astronomical Telegrams

3.1. Weather

The weather category reports the fraction of scheduled observing time that weather conditions are suitable for observing. Table 3.1 summarizes the amount of time lost to weather and Figure 3.1 plots the fraction of suitable observing time against the baseline forecast. Averaged over the quarter, the fraction of available observing time was greater than predicted in the baseline plan, which led to solid gains in data acquisition for all three surveys.

Table 3.1. 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	449	146	68%	60%
Dark & Gray Time	696	188	73%	60%

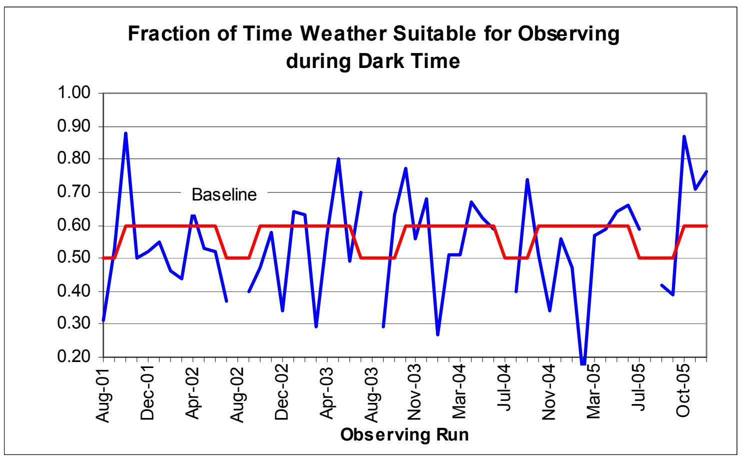

Figure 3.1. Percentage of Time Weather Suitable for Observing

3.2. System Uptime

System uptime measures the availability of equipment when conditions are suitable for observing. We averaged 95% uptime against a baseline goal of 90%. Table 3.2 summarizes the total amount of time lost to equipment or system problems and Figure 3.2 plots uptime against the baseline goal. We lost observing time to several minor problems that occurred at various times throughout the reporting period.

Table 3.2. Potential Observing Hours Lost to Problems in Q4

Observing Condition	Total hours potentially available for observing	Total hours lost to problems	System Uptime	Baseline Forecast
Dark Time	449	20	96%	90%
Dark & Gray Time	696	35	95%	90%

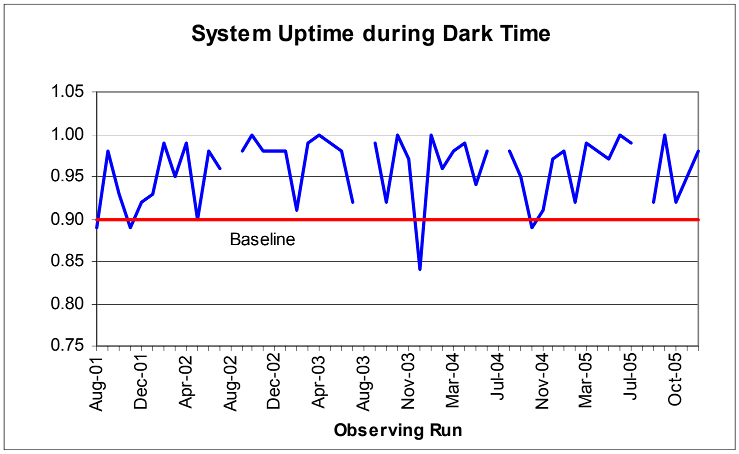

Figure 3.2. System Uptime

3.3. Imaging Efficiency

Imaging efficiency averaged 90% against a baseline goal of 86%. Favorable weather allowed long imaging runs on Stripe 82 for the Supernova Survey. As imaging overhead time is the same regardless of run length, longer runs improve efficiency.

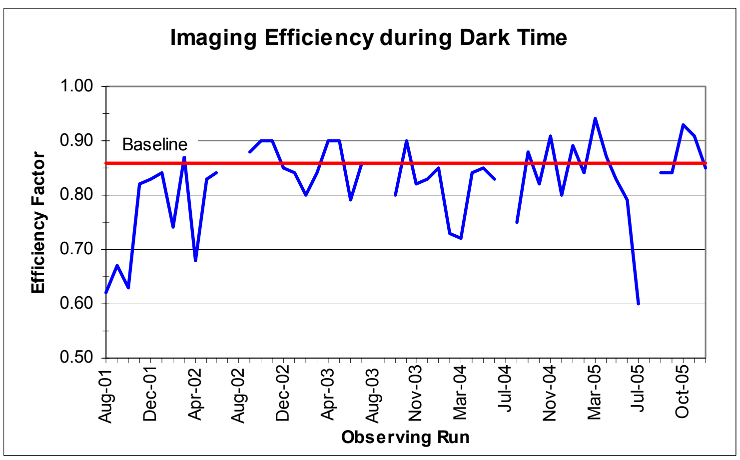

Figure 3.3. Imaging Efficiency

3.4. Spectroscopic Efficiency

Spectroscopic efficiency is derived by assessing the time spent performing various activities associated with spectroscopic operations. Table 3.3 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, spectroscopic efficiency averaged 65% against the baseline goal of 64%.

Table 3.3. Median Time for Spectroscopic Observing Activities

Category	*Baseline*	*Run starting Sep 21*	*Run starting Oct 23*	*Run starting Nov 19*
Instrument change	10	0	8	6
Setup	10	18	7	10
Calibration	5	5	5	5
CCD readout	0	3	3	3
Total overhead	25	26	23	24
Science exposure (assumed)	45	45	45	45
Total time per plate	70	71	68	69
Efficiency	0.64	0.63	0.66	0.65

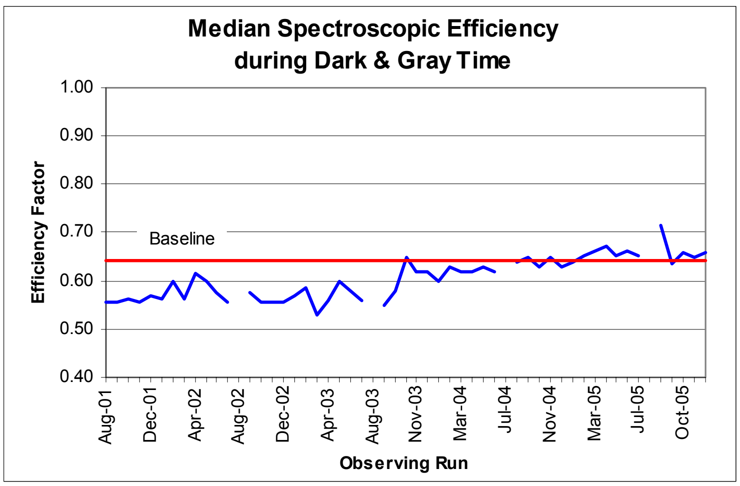

Figure 3.4. Spectroscopic Efficiency

4. OBSERVING SYSTEMS

Observing systems includes the instruments, telescopes, computers and various sub-systems that support observing operations at APO.

4.1. The Instruments

The imaging camera saw extensive use in the quarter, as a significant amount of observing time was devoted to imaging for the Supernova Survey. In fact, 84% of the available observing time in the October run was devoted to imaging. The camera performed well for the most part, with the exception of a brief period between September 30 and October 11 in which the i5 chip became problematic. A number of non-intrusive measures were taken to diagnose the problem (e.g., power cycling and adjusting CCD voltages) but none had any effect. However, on October 11 the chip started behaving normally again. As of this writing, we have not seen the problem recur, nor do we understand what caused the problem in the first place. It is worth noting that we did not lose science time because of this problem, as it was determined that the impact of the absence of the i5 chip on the supernova data was tolerable.

We experienced problems with the u3 chip on one night during the November run. The chip returned to normal after four hours in the enclosure; we speculate the problem was temperature-related, as there was a 3.5°C temperature difference between the time in which the problem was present and when it was not.

The spectrographs performed reasonably well during the quarter. We had a minor problem with fogging in the lens assembly of the blue camera (b1) on spectrograph 1 (sp1). During the summer shutdown, the spectrographs were removed from the telescope and the cameras stored in the plug-plate lab for safe keeping. Humidity levels are high at the site during the summer monsoon season and the plug-plate lab is not air-conditioned. We speculate that moisture accumulated in the lens assembly during the shutdown and condensed during the fall as temperatures began to decrease. An inspection of the lens assembly showed that the interior lens was fogged, but not the dewar lens. This was fortunate, as the dewar lens is composed of calcium fluoride (CaF_2), which is hygroscopic and subject to damage by moisture. We removed the moisture from the interior lens by applying a dry air purge to the lens assembly. As a precautionary measure, we also heavily purged the red camera lens assembly, as well as the lens assemblies for both cameras on spectrograph 2. After purging for several days, the lens assembly was reinstalled on the telescope and bias data taken to verify that the fogging was gone and that spectrograph throughput was as expected. We did not lose science time to this problem, as it occurred during the time allocated to the Supernova Survey, when imaging is the primary mode of operation.

To preclude moisture problems in the future, a permanent purge system was attached to the four camera lens assemblies. The lens assemblies had been machined to accommodate purge fittings, but fittings had never been installed. The engineering team installed fittings into the lens assemblies and connected them to the site dry air system. The spectrograph lens assemblies are now under constant purge.

We experienced a problem with the liquid nitrogen (LN_2) autofill system on the spectrographs. Troubleshooting indicated that the problem was associated with the fill sensors in the intermediate LN_2 dewars. Adjusting the fill set points to accommodate for drift in the sensors solved the problem. In the process of troubleshooting the fill problem, however, we also discovered that the power supply for the autofill system was only marginally capable of meeting the demands of the system. We upgraded the power supply to a higher capacity unit.

In the process of troubleshooting the autofill system, we discovered that the valves in the LN_2 lines have removable valve seats and that the valves can be purchased with different orifice sizes. Upon further inspection, we found that some of the valves in the system had different size orifices. Replacements were made so that all valves now have the same size orifice, which will better balance flow rates and pressures in the system.

We had a glue failure between one of the fiber bundles and the slithead terminal block on fiber cartridge #1. This cartridge has been in service the longest and this was the first such adhesive failure we have experienced. The fiber bundle was re-glued to the terminal block and the alignment verified; it took two attempts to get the alignment correct. During the process, a broken fiber was discovered in the adjacent fiber bundle. Modifications were made to the mapping software to accommodate the broken fiber and re-aligned bundle. As we have nine cartridge assemblies, no science time was lost while these repairs were made.

4.2. The 2.5m Telescope

Supernova observing operations in Q4 extended into the moony periods between dark runs, which is typically the time scheduled to perform routine engineering and maintenance activities on the telescopes and instruments. With the time to perform routine work limited, engineering work in Q4 instead focused on addressing problems as they arose, and on activities away from the telescopes and instruments.

The following list highlights some of the engineering work performed during the quarter.

1) We finished the fabrication of a spare M2 Galil controller assembly. The new unit was extensively bench-tested, installed and tested on the telescope, and certified as a working spare. We also modified the M2 Galil controller assemblies so they can not be mistakenly installed into the M1 Galil controller enclosure.

2) During the summer shutdown, a new pump motor was installed on the site ground loop cooling system. In Q4, we installed a new heat exchanger onto the pump housing to reduce the thermal load into the telescope enclosure.

3) In the spring of 2005, we fabricated a prototype standalone timer circuit that could be triggered by the imager's LN2 autofill system. The standalone timer circuit was used to temporarily bypass a faulty imager dewar fill sensor. The prototype worked well, so we constructed a more robust version in the event that we have further problems with the imager fill sensors. Long-term, we plan to fabricate and permanently mount two of these units onto the imager's saddle, to minimize downtime should a sensor fail during observing operations.

4) A new enclosure control pendant was installed on the west end of the 2.5m telescope enclosure.

5) We started testing prototype argon/neon calibration lamps, as the units we have been using are no longer available. We anticipate certifying new lamps by early spring.

6) In the process of training the observers in the use of the emergency closing generator, we discovered that the generator was unable to supply sufficient power to move the enclosure. We found that gasoline had seeped past the piston rings and into the crankcase of the motor, which drastically reduced the output power of the engine. We drained and replaced the oil from the crankcase and replaced fouled spark plugs; output power returned to normal and the generator is once again capable of moving the enclosure. A stop valve was installed in the fuel line to preclude flooding during long periods of non-use. We also modified our monthly generator testing procedure to include moving the enclosure, as opposed to simply running the generator.

7) We continue to have problems with the drive system on the DIMM telescope. The problem stems from the fact that our duty cycle is much heavier than that for which the telescope was designed. We are in the process of evaluating replacement telescopes. We envision purchasing and installing a new 10-inch class telescope in 2006-Q1.

8) We discovered a slight misalignment in the contact line between the rotator motor drive capstan and the rotator drive surface. We re-aligned the capstan to preclude further damage in the near-term. We plan to re-grind the capstan surface during the 2006 summer shut down.

4.3. The Photometric Telescope

The Photometric Telescope (PT) worked reasonably well throughout the quarter. We continued to troubleshoot occasional communication problems with the Telescope Control System (TCS), as discussed in the last report. Running out of ideas, we put into service the spare TCS computer; the communication problems appear to have gone away. We are in the process of replacing the computer (the control software is DOS-based and requires an 80486-class machine). We are also in the process of procuring a new controller card from the vendor, to deepen our spares pool and avoid a costly full upgrade of the system.

In December, the transfer lines from the Cryotiger closed-cycle refrigerator became jammed in the PT right ascension axis fork tines. The SEGUE Survey requires observing in new areas of the sky, which caused the Cryotiger lines to move into the fork tines. Fortunately, no permanent damage was done. To preclude further problems, we rerouted the lines by fabricating a better mounting system. We also plan to install range-of-motion limit switches for the right ascension axis.

4.4. Operations Software and the Data Acquisition System

The Telescope Control Computer (TCC) software was upgraded with a new version of Multinet, the software used by the TCC for network communications such as TCP/IP. The upgrade was necessary to implement a version of the NTP server that the new DA system can access to properly obtain a time signal. Work began on TCC v2_9, which will improve pointing near the zenith by computing the effect of field rotation in the pointing model. We anticipate testing and implementing the upgrade in 2006-Q1.

The Telescope Performance Monitor (TPM) was upgraded to accommodate the port from the IRIX operating system to Linux on the new host computer, *sdsshost2*. In particular, the TPM displays were built with an old and now unsupported display manager called "dm". The old display manager was not developed under Linux, which caused problems when the TPM displays were moved to *sdsshost2*. We installed a new display manager, “edm”, on sdsshost2 and moved the TPM displays to run under it.

We continued commissioning and debugging the new DA system, and porting code from *sdsshost* to *sdsshost2*. As anticipated, we experienced a number of minor problems during the porting and commissioning process. With solid developer support and the implementation of temporary workarounds, none of the problems resulted in lost science time.

As previously reported, the DA upgrade eliminated the process of writing data to DLT tapes. Data are now written to disk and transferred to Fermilab and Princeton for processing via the Internet. Initially, the bandwidth off the mountain (~12.9 Mbits/sec) was just adequate to keep up with daily data acquisition rates. It took approximately two hours to transfer one hour’s worth of data so barring problems, a night’s worth of data would be transferred off the mountain before the next night of data acquisition began. To increase this buffer and decrease transfer times, we contracted with the ISP to increase the microwave bandwidth to approximately 18 Mbits/sec. As a result, data transmission rates were measured to be ~40% faster, with maximum sustained rates measured at 17.8 Mbits/sec. Figures 4.1 and 4.2 show data transfer rates measured before and after the microwave link upgrade.

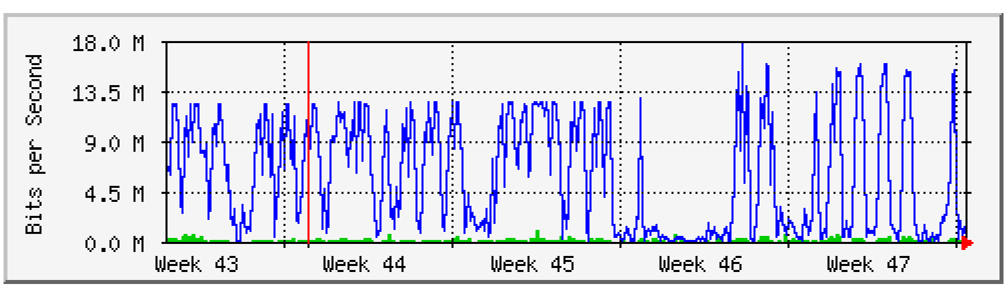

Figure 4.1. Data transfer rate from APO prior to the microwave link upgrade.

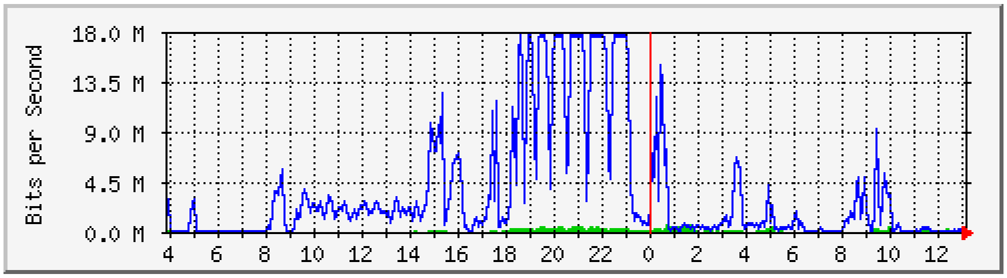

Figure 4.2. Data transfer rate from APO measured starting on 11-27-05, after the microwave link upgrade was in place.

Ongoing work associated with the DA upgrade includes fixing bugs as they occur to improve system performance, reliability, and robustness. Unless they are critical-high in nature and impede our ability to collect data, bug fixes are only being implemented during “shake” periods, following standard observing software change control procedures. We are using the SDSS Problem-Reporting Database, Gnats, to record and track problems and fixes. We are working to implement the process of writing a backup copy of all data to removable disks at APO, and refining and formalizing the data transfer process between APO and production data processing operations. Finally, there are a few programs remaining on *sdsshost* that need to be ported to *sdsshost2*. Once the port is complete, we plan to de-commission *sdsshost* and eliminate the high maintenance contract costs.

5. DATA PROCESSING AND DISTRIBUTION

5.1. Data Processing

5.1.1. Pipeline Development and Testing

No changes were made to the production Legacy photometric or spectroscopic pipelines in Q4.

Work continues at Princeton on the spectroscopic pipeline, idlspec2d, v5. The pipeline has been set up to run robotically and nearly all Legacy plates have been run through the development pipeline. In time, SEGUE plates will also be processed through the pipeline. Software testing and debugging is ongoing; the test version is currently a bug-fix variant of idlspec2d, v5_1. A testing plan is being developed to compare the differences between the outputs of the current version of the spectro 1d pipeline and idlspec2d v5_1. Scripts are being developed to automate the testing process.

Work continues by the JINA-MSU team on the development of the stellar atmosphere pipeline that will be used for SEGUE observations in order to estimate atmospheric parameters (Teff, log g, [Fe/H]) based on R = 2000 spectroscopy and ugriz photometry.

The stellar atmosphere pipeline in use is presently based on a number of independent methods (obtained from different calibrations) for each parameter, which are then suitably averaged in the final estimation process. Estimates of the internal scatter in the determination of a given parameter are also kept track of. The pieces of this pipeline are (for now) run as separate procedures. Steps are being taken in order to put all of these pieces into an IDL procedure so that they can be run in standalone mode. Several of the calibrations rely on input B-V colors, which have to be derived from observed g-r colors via an approximate transformation. This step will be eliminated soon, as new calibrations tied directly to ugriz colors are obtained. Highlights of Q4 progress in this area are as follows:

- Generation of a full grid of synthetic spectra for carbon-enhanced versions of the MARCS models, covering a wide range of [C/Fe]. This is an extension on the sub-grid that was generated in the previous quarter

- Application of the V1.0 line index definitions (and assigned errors, based on the quality of the spectra), as required in the original agreement. These errors have been checked, and appear to be quite reasonable, relative to expectation. It is anticipated that further refinement of these line indices will take place as SEGUE evolves

- Development of an "optimal" continuum fitting procedure (in IDL), which has been demonstrated to be superior to previously available methods we were using

- Continued testing of methodologies (based on Artificial Neural Networks and Support Vector Machines) for the estimation of stellar atmospheric parameters based on ugriz photometry alone. It is expected that these techniques might prove useful for refinement of SEGUE target selection algorithms, or as standalone methods for further analysis of stars without available spectroscopy. In addition, we have initiated further testing of two additional methods, based on Self Organizing Maps, and the CLASSX algorithm.

- Continued obtaining high-resolution spectroscopy of SDSS stars with predicted parameters from the present pipeline, so that calibration and refinement of these estimates can be carried out. To date, some 20-30 stars have available Keck spectroscopy, and of greatest importance, over 50 HET spectra have been obtained in the first (of three) trimesters anticipated to be obtained over the course of the coming year.

In the coming quarter, our primary goal is to have a working (IDL-based) spectroscopic pipeline up and running in time for the March SDSS-II Collaboration meeting. This will represent V1.0 of the pipeline. Once this goal is accomplished we will seek to refine and improve the pipeline based on empirical comparison of the estimated atmospheric parameters with those obtained from high-resolution spectroscopic observations that we either have already, or are presently obtaining. This effort will also provide an external measure of the expected errors in the parameter determinations. Note that an AAS paper (Sivarani et al.) at the Washington DC AAS meeting presented a first-pass effort on this (based on ~ 10 stars).

We plan to run V1.0 of the pipeline through all available SEGUE data as of mid-February, or early March, in order to verify that it is working properly (i.e., that it reproduces what we obtain from the "individual components" of the pipeline mentioned above). A catalog of the derived parameters will be circulated to the SDSS-II/SEGUE collaboration once this is done. Furthermore, the intention is to have a discussion of the present version of the calibration comparison of pipeline results at the March Collaboration meeting.

5.1.2. Data Processing Operations at APO

The new SN Compute Cluster at APO was used to process all data obtained for the SN Survey. New data were typically processed within 24 hours of acquisition in order to quickly identify Type Ia supernova candidates for spectroscopic follow-up. Data processing at APO was performed by members of the SN team. Future work on the compute cluster will include minor modifications as necessary to fix bugs, fine-tune performance and increase operational efficiency and system robustness.

5.1.3. Data Processing Operations at Fermilab

Data for the Legacy, SEGUE and SN Surveys are processed at Fermilab. In Q4, we processed spectroscopic data from 28 Legacy plates and 17 SEGUE plates. All spectro data were processed in the standard production environment using existing Legacy versions of the spectroscopic pipelines (idlspec2d VV4_10_4 and spectro 1D v5_9_4) and flat fields from 2004-2005 observing season.

In addition to spectro processing, newly acquired imaging data were processed at Fermilab. All imaging data were processed in the standard production environment with the Legacy version of the photometric pipeline (photo v5_4_28). Because some of the SN data were obtained under less-than-ideal conditions (i.e., non-photometric, moony, and/or cloudy), we are unable to calibrate those data using standard pipelines. As time permits, we will attempt to calibrate those data taken under photometric conditions.

We experienced a number of problems with computer hardware associated with data processing operations. We experienced disk drive and controller problems on several of the file servers used to support data processing operations; the problems resulted in file corruption and required either the retrieval of data from the Enstore tape robot, or the re-generation of data files. In mid-December, a disk failed on the file server hosting the project's software code repository (CVS) and software products. The faulty disk was replaced and the system rebooted. The CVS repository is backed up nightly, so the repository was restored from the previous night's back-up. Changes checked in to CVS between the time of the backup and disk failure were lost; fortunately, only a small number of changes had been checked in during this time. The products directory was restored from a two-week-old backup and then additional changes made and products cut to replace the work lost between the time of the last backup and the disk failure. Approximately three working days were lost to the recovery effort.

A significant amount of effort went into developing a robust system for automatically transferring new data from APO, packaging it for processing at Fermilab, and copying it into the Enstore tape robot for long-term archiving. A scripted process was developed and implemented that automatically searches the APO file server for new data, initiates a data transfer process when new data is discovered, compares checksums to verify the integrity of the transferred data, and verifies that all data associated with a particular imaging scan or plate exposure exists and is properly accounted for. Warning messages are issued when problems are identified. A web page was also developed and implemented to record the status of data transfers and to inform the observers of those data that were successfully transferred and can be safely removed from the fileserver at APO.

5.1.4. Data Processing Operations at Princeton

Legacy and SEGUE data are also being processed at Princeton. New computer hardware was acquired and put into service to support operations. The reduction environment for imaging data was set up and data were run through Photo in preparation for running *ubercal*. All Apache Wheel imaging scans have been reduced, as well as several of the previously unreduced regular survey runs. Additional work related to the calibration effort included testing *ubercal*, testing the reproducibility of Photo on different compilers, and training new personnel in how to run the data reduction code.

The reduction environment for spectroscopic data was set up to run in robotic mode. Spectroscopic data from the Legacy and SEGUE Surveys are being processed through idlspec2d v4_10_7, while development work continues on idlspec2d v5.

Data transfers from APO and Fermilab are now fully automated. Spectro data is transferred from APO and imaging data is transferred from Fermilab; both transfers occur via the internet. The process to transfer data from APO has been in place for some time. The process of transferring imaging data from Fermilab was newly implemented, as this is a new mode of operation made possible by the DA upgrade.

5.2. Data Distribution

5.2.1. Data Usage Statistics

To date, the general public and astronomy community have access to the EDR, DR1, DR2, DR3, and DR4 through the DAS and SkyServer interfaces.

In past reports we included graphs showing data volume transferred from the DAS via rsync and through a web interface. Due to problems with the DAS in Q4 (discussed below), we were unable to accurately log these usage statistics. Logging has recommenced in January and will be reported in the 2006Q1 report.

Figure 5.1 plots the number of web hits we receive per month through the various SkyServer interfaces. In Q4 we recorded 16.4 million hits, compared to 15.6 million hits in Q3 and 12.7 million hits in Q2.

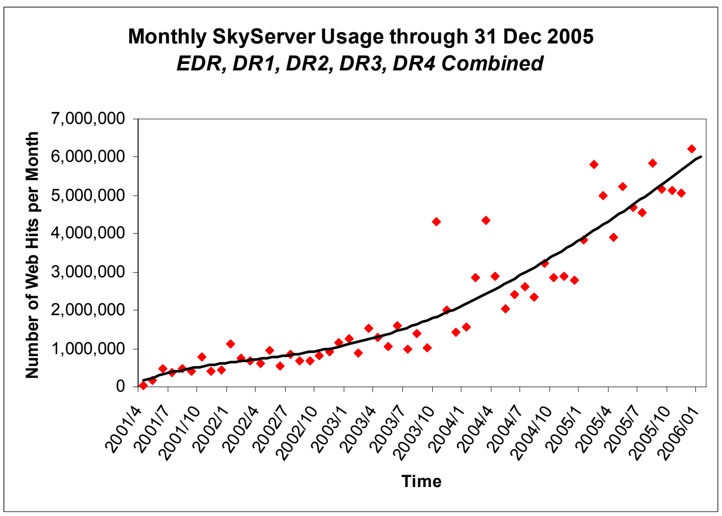

Figure 5.1. SkyServer usage per month, for all public releases combined.

Figure 5.2 shows the total number of SQL queries executed per month. We executed 1.7 million queries in Q4, compared to 0.7 million queries in Q3, and 1.6 million queries in Q2. As previously reported, the spike in February was due to machine-generated queries.

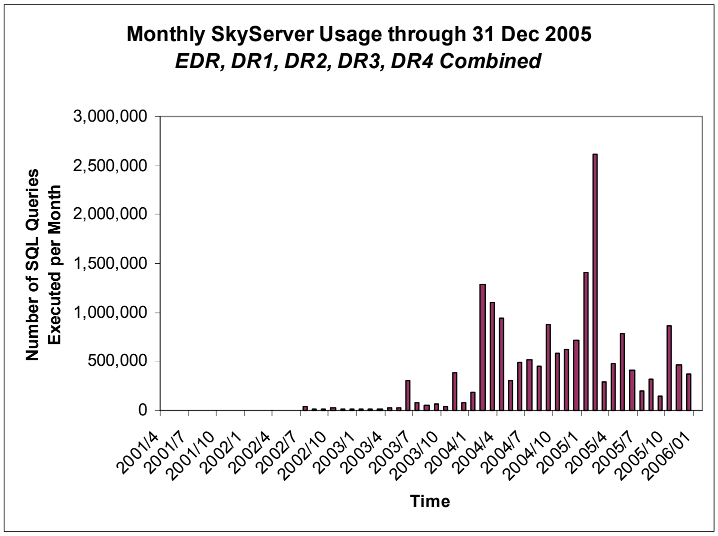

Figure 5.4. SkyServer usage, measured by the number of SQL queries submitted per month.

Through December 31, 2005, the SkyServer interfaces have received a total of 125 million web hits and processed over 17.6 million SQL queries. As the graph shows, we are experiencing a phenomenal usage rate with usage continuing to grow steadily.

5.2.2. Data Release 5

The initial version of Data Release 5 (DR5) was made available to the collaboration on October 21, 2005. DR5 contains all survey quality imaging data collected through June 30, 2005, and the corresponding spectra. Thus, DR5 comprises the final Science Archive of the SDSS-I. Table 5.1 summarizes the contents of DR5 and provides a comparison with the DR4 release. The table and following summary of DR5 contents is reiterated for convenience from the 2005-Q3 report.

Table 5.1. DR5 Contents

	DR4	*DR5*	*Increment*
Imaging			
Footprint Area	6,670 sq. deg.	8,000 sq. deg	1,330 sq. deg.
Imaging Catalog	180 million objects	215 million objects	39 million
Data volume			
Images	7.5 TB	9.0 TB	1.5 TB
Catalogs (DAS, fits format)	1.5 TB	1.8 TB	0.3 TB
Catalogs (CAS, SQL database)	3.0 TB	3.6 TB	0.6 TB
Spectroscopy			
Spectroscopic Area	5,320 sq. deg	5,740 sq. deg.	420 sq. deg.
Total Number of Spectra	849,920	1,048,960	199,040
Galaxies	565,715	674,749	109,034
Quasars (redshift < 2.3)	67,382	79,934	12,552
Quasars (redshift > 2.3)	9,101	11,217	2,116
Stars	102,714	154,925	52,211
M Stars and later	50,373	60,808	10,435
Sky Spectra	44,363	55,555	11,192
Unknown	10,272	12,312	2,040

In addition to standard survey data, DR5 contains spectroscopic data from 361 “extra” and “special” plates (compared to 276 for DR4; Δ = 85 plates). The bonus data come from the following plate observations:

- 62 "extra" plate/MJD combinations which are repeat observations of 53 distinct main survey plates (*increment of one plate above that released with DR4*).
- 289 distinct "special" plates, which are observations of spectroscopic targets, mostly in the southern galactic cap, which were selected by the collaboration for a series of specialized science programs. Some of these plates are outside of the DR5 imaging area (*increment of 83 plates above that released with DR4*).
- 10 "extra-special" plates, which were repeat observations of "special" plates (*increment of 1 plate above that released with DR4*).

Efforts in the early part of the quarter focused on getting the DR5 data set to the collaboration. Once DR5 was released, efforts focused on tying up a few loose ends and adding a small number of enhancements that will be included in the final data release version. Details are provided in Section 5.2.4 below. We anticipate releasing the “enhanced” and final version of DR5 to the collaboration by the end of March 2006, and to the public on June 26, 2006, in accordance with the approved data distribution plan.

5.2.3. Data Archive Server

Work on the Data Archive Server (DAS) in the early part of the quarter focused on final preparations for the DR5 collaboration release. The creation and testing of links to the DR5 data set was completed and the DAS DR5 web pages created.

On November 22, two of our Linux web servers were compromised by a computer hacker. In addition to hosting static web pages associated with data processing operations, these machines served the DAS to the collaboration and general public. The compromised servers were quickly taken offline and secured. While the compromised servers were being investigated, replacement machines were brought into service and the process of restoring the static web pages and rebuilding the DAS was started. Restoring the static pages occurred within a few days but rebuilding the DAS took substantially longer. The underlying DAS infrastructure had developed over time in a somewhat ad-hoc manner. We viewed the rebuild as an opportunity to clean out artifacts and unused code from earlier data releases and implement a more robust system that would be easier to expand and maintain going forward. Restoring the DAS required rebuilding all of the DAS links for all data releases, which proved to be a time-consuming process. The DR4 links were restored and brought online on December 12. The links for DR3, DR2 and DR1 were restored and brought online near the end of December. At quarter’s end, we were working on restoring the DR5 links.

The DAS downtime obviously affected our usage statistics. No data were transferred after November 22. Prior to November 22, we had problems with our logging process. The harvesting process for data transfers via the web interface stopped on September 30. The harvesting process for rsync transfers stopped logging transfers between October 1 and October 30. Harvesting of rsync transfer volumes started again on November 1 and ended on November 22, when the servers were taken offline. In the past, we were able to recover these logs manually. In this case, however, the information was lost because the harvesting data was on the machines that were compromised and secured. We plan to implement a more robust harvesting process on the new servers in order to more reliably track and report usage statistics.

5.2.4. Catalog Archive Server

Work on the Catalog Archive Server (CAS) included finishing up the DR5 data load for the October collaboration release; and testing and debugging a new version of the CAS Finish step. The new Finish step includes the region, wedge, and sectors code that will allow collaboration members and other users to generate large scale structure galaxy sample catalogs, including mask and selection function information, directly from the database. During the quarter, the enhanced Finish step was installed at Fermilab and run on a production copy of BestDR5 by the production group. Backup copies were made of the enhanced BestDR5. At quarter's end, the development team was running tests against the enhanced BestDR5 to validate the new outputs, and developing sample queries and documentation to help users take advantage of the new features.

An enhanced version of the DR5-CAS will be released to the collaboration in 2006-Q1. Remaining work prior to release includes finishing the documentation and sample queries for the region/wedge/sector code and adding RunQA data, which is generated during data processing, to tables in BestDR5 and TargDR5. In addition, two separate groups within the collaboration have computed photometric redshifts (photo-z's) using different algorithms. Both groups are in the process of providing tables containing computed photo-z outputs, along with estimated errors, to the production team for loading into BestDR5. In addition to providing data outputs, the two groups are providing documentation describing the algorithms used to compute the photo-z's.

In addition to work on the DR5-CAS, a significant amount of effort went into loading data into the Runs database (RunsDB). Our intent is to load into the RunsDB, all of the imaging runs collected since the start of SDSS-I operations and all of the imaging runs collected for SDSS-II as they become available. In Q4 we loaded nearly all of the Stripe 82 runs into the RunsDB; these are the repeat imaging scans on the southern equatorial stripe. We loaded the data in two steps (an initial load of 41 runs, followed by a Finish step, followed by a follow-on load of 20 runs, followed by a final Finish step). The initial load was tested and then made available to the collaboration on December 6. At quarter's end, a few members of the collaboration were running sanity tests against the full Stripe 82 load, in preparation for a collaboration release in early January. Once the Stripe 82 runs are released, we will begin loading SEGUE imaging scans into the database. Based on our experience loading the Stripe 82 scans, we expect to release the RunsDB with the SEGUE incremental load in mid-to-late February. We will follow the SEGUE load by loading SDSS-I imaging scans on a "year-by-year" basis; that is, we will load imaging data collected in the 2004-2005 observing season, followed by data collected in the 2003-2004 season, etc. At the end of each "season load", we will make a backup copy of the database and then make that version available to the collaboration for use, while production loading continues on the master copy. We anticipate releasing updated versions at roughly six-to-eight-week intervals until all of the imaging data are loaded.

In addition to supporting production operations, the CAS development team worked on activities designed to improve performance and maintainability going forward. The sizes of the Best and Target databases are getting quite large (2.3 TB for BestDR5; 2.0 TB for TargDR5). To improve performance and make it easier to move these databases around, we intend to partition the databases into several smaller volumes. Scripts are under development that will automatically perform the partitioning process. BestDR4 has been partitioned and tests are underway to measure and quantify performance improvements. Based on the success of early results, we anticipate implementing a partioning scheme on the final versions of BestDR5 and TargDR5, prior to the 2006-Q1 collaboration release described above.

We are currently deploying our CAS databases on Microsoft SQL Server 2000. The development team has been porting BestDR4 to SQL Server 2005 at JHU; the new version has many features designed to improve performance on very large databases. In addition to addressing known bugs, implementing SQL

Server 2005 will allow us to take advantage of features that will allow us to simplify some of our current indexing schemes. Barring unexpected problems, we anticipate loading DR6 using SQL Server 2005.

Upgrades were made to the CasJobs web service in Q4. New features, implemented in CasJobs v2_8_3 include a query plan button that shows the user the estimated graphical query execution plan; improved syntax checking; and VOPlot plotting capability. All changes were checked into the CVS code repository. In addition to code changes, we set up a CasJobsTest site onto the production servers at Fermilab. The addition of the test site provides the developers with the infrastructure to more thoroughly test and debug changes before releasing the changes onto the production machines, as further described below.

We implemented several infrastructure upgrades to our production systems in Q4. We installed and deployed a three-node cluster of new web servers on the front end of the CAS production databases. The web servers were installed in a load-balancing configuration to better distribute the user load coming into the system. We then redeployed two of the existing web servers in a "development and test configuration." One of the web servers is designated for development use; the second for testing. Under our new process, code and system-level changes are developed and debugged on the development machine. Once changes are deemed "production-ready", they are pushed to the "test" machine for final testing. If testing is successful and no further changes are necessary, the code is pushed to the production machines in a scheduled manner. If testing indicates that further changes are required, the changes are made on the "development' machine and the testing and evaluation process starts over. The new hardware configuration and procedure formalizes our development and testing process, which we anticipate will lead to more robust deployment of code into the production operation.

We initiated the procurement of three new file servers to support the Runs DB loading operation and to replace two existing machines nearing the end of their useful life. Near the end of the quarter, we started working on the development of online performance monitoring tools to track and log measures such as CPU usage, available disk space, available memory, etc. Gathering and analyzing such data will help us assess how well our systems are performing and identify areas in which performance needs to be tuned. These tools will also allow us to implement alerts to warn of pending problems before they affect system performance and availability. We anticipate implementing a prototype performance monitoring system by the end of 2006-Q1.

6. SURVEY PLANNING

6.1. Observing Aids

Several programs are used to aid in planning and carrying out observations. No software changes were made. Additional patches for low-priority stripes in the SEGUE program were loaded into the patch database. The plate inventory database was modified to allow the observers to more easily identify the short and long exposure plates of a SEGUE tile so both are more likely to be observed within the same dark run.

6.2. Target Selection

For this quarter, 74 plates were designed and drilled in two drilling runs. Of these, 36 were for the Northern survey area, 19 were for the normal exposure SEGUE plates, 17 were for double length exposure SEGUE plates, and two were special SEGUE cluster plates.

6.3. 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 have been made.

7. COST REPORT

The operating budget that the Advisory Council approved for SDSS-II activities during the period July 1 through December 31, 2005 consists of $355K of anticipated in-kind contributions from Fermilab, the US Naval Observatory (USNO), the University of Chicago (UC), the Johns Hopkins University (JHU), the University of Washington (UW), and the Joint Institute for Nuclear Astrophysics (JINA); and $2,305K for ARC-funded cash expenses.

Table 7.1 shows actual cost performance for ARC-funded cash expenses in Q4. A more complete table comparing actual to baseline performance is included in the appendices of this report. Appendix 1 compares Q4 cash expenses to the budget and presents the revised cash forecast for 2005. Appendix 2 compares actual in-kind contributions to the budget and presents the revised in-kind forecast for 2005.

Table 7.1. Q4 Cash Expenses and Forecast for 2005 ($K)

	2005 – 4th Quarter		2005 Operations Budget Total *(for the period Jul-Dec 2005)*	
Category	Baseline Budget	Actual Expenses	Baseline Budget	Actual Expenses
1. Survey Management	61	56	208	168
2. Survey Operations				
2.1. Observing Systems	187	190	376	356
2.2. Observatory Support	394	438	787	805
2.3. Data Processing	118	116	236	199
2.4. Data Distribution	24	28	47	28
2.5. ARC Support for Survey Ops	6	0	11	0
3. New Development				
3.1. SEGUE Development	33	58	138	95
3.2. Supernova Development	0	0	162	74
3.3. DA Upgrade	30	39	241	256
3.4. Photometric Calibration	5	5	10	10
4. ARC Corporate Support	13	13	32	16
Sub-total	869	914	2,250	1,979
5. Management Reserve	28	0	55	0
Total	897	914	2,305	1,979

7.1. Q4 Performance - In-kind Contributions

The sum of in-kind contributions in Q4 was $217K against the baseline forecast of $122K and was provided by Fermilab, JHU, UC, USNO, UW, and Michigan State University (MSU) for JINA, as follows:

- Fermilab provided support for survey management, data processing and data distribution activities. Effort was also provided to support oversight and planning, and development work for the SEGUE and Supernova projects. The level of effort provided to support development, data processing and data distribution was greater than anticipated, as were the salary costs of the individuals performing this work.
- JHU provided support for the development, loading and hosting of the databases associated with the CAS, CasJobs, and SkyServer, at the anticipated level.
- UC provided support for the development of a program to monitor the performance of the on-mountain SN data system, checking detection efficiency and photometric accuracy as functions of magnitude for each run. UC also provided support for the hand scanning of supernova candidates. The level of effort provided was as anticipated.
- USNO provided support as required for the astrometric pipeline and other software systems they maintain. The level of effort required was less than anticipated in the baseline.
- UW contributed the overhead associated with the plate drilling operation as anticipated.
- MSU provided support for the development of the spectroscopic pipelines that will be used for SEGUE observations in order to estimate atmospheric parameters. The level of effort provided was as anticipated.

The value of in-kind contributions in Q4 exceeded the baseline forecast because the value of the Fermilab contribution was greater than anticipated, and because the baseline underestimated the level of anticipated Q4 in-kind contributions. When the baseline was prepared, the distribution of in-kind contributions was too heavily front-loaded into Q3; the baseline forecast overestimated the level of effort for development in Q3. Since the total level of anticipated support for the two quarters was fixed, overestimating levels of effort in Q3 had the effect of underestimating levels in Q4. As a result, actual contributions in Q4 exceed the baseline forecast. Notwithstanding, the level of effort provided matched that required to complete the required work.

7.2. Q4 Performance – ARC Funded Cash Expenses

ARC-funded expenses were $914K, or $30K (3%) above the fourth quarter budget of $884K, excluding management reserve.

Survey management costs were $56K against a budget of $61K. Actual support costs for the Director, Project Scientist, Public Information Officer, and project management support staff were less than anticipated. No charges were incurred against the budget for Public Affairs or Collaboration Affairs. For the year, the revised forecast for Survey Management expenses is $168K, or $40K (20%) below the baseline budget.

Observing Systems costs were $190K against a budget of $187K. UW costs were less than budgeted, as the amount of UW engineering and technical effort required to support on-going operations was less than anticipated. Actual Princeton and Fermilab expenses exceeded the budget. The baseline Princeton budget did not include salary support for technical staff who participated in the summer maintenance work on the imaging camera. The baseline Fermilab budget provided for 50% of the salary support for the Telescope Engineer; in reality, the project provided for 75% of the salary support for this individual. The overrun in these two areas is covered by re-allocating funds for additional engineering support and unanticipated hardware expenses that are held in the ARC account for Observing Systems Support. Direct expenditures from the ARC account in Q4 were associated with the upgrade of the Uninterruptable Power Supplies for the APO computer room; the budgeted upgrade was implemented to offset the added heat load imposed on the computer room from the Supernova Compute Cluster. For the year, the revised forecast for Observing Systems expenses is $356K, or $20K (5%) below the baseline budget of $376K.

Observatory Support costs were $438K against a budget of $394K. Salary expenses were below budget, largely due to the flux in filling open personnel positions. Travel expenses were also below budget. Operations costs were greater than expected, in part due to higher expenses for the increased-bandwidth internet and higher than normal utility and fuel costs. Operating costs were also higher than expected due to non-recurring staff realignment and recruitment expenses. Also, there were outstanding encumbrances for annual maintenance contracts that were closed out in Q4; these contracts are renewed on an annual basis and planned for in the annual operating budget. For the year, the revised forecast for Observatory Support expenses is $805K, or $18K (2%) above the baseline budget of $787K.

Data Processing costs were $116K against a budget of $118K. Incurred costs at Fermilab, Princeton and Chicago in Q4 were in close agreement with the budget. For the year, the revised cost forecast for Data Processing is $199K, or $37K (16%) below the baseline budget.

Data distribution costs were $28K against a budget of $24K. Fermilab expenses exceeded the budget slightly because we added staff to better support data distribution operations. For the year, the cost forecast for Data Distribution is $28K or $19K (63%) less than the baseline budget of $47K. The variation is due to the fact that Q3 expenses associating with finishing the DR5 load were charged against the SDSS-I budget, not the SDSS-II budget.

No expenses were incurred in Q4 against the ARC accounts holding funds for additional Survey Operations support (specifically, the accounts for Additional Scientific Support and Observers' Research Support). As no expenses were incurred in Q3 either, actual expenses for the year are zero against a budget of $11K.

Expenses associated with development work for the SEGUE Survey were $58K against a budget of $33K. Cash-funded development work on SEGUE occurred at Princeton and Fermilab. When the baseline budget for the Princeton effort was prepared, it was anticipated that development work would occur in Q3 and so the budget was front-loaded. In reality, resources at Princeton were not available as early as anticipated so the work was spread over Q3 and Q4. The total budget for the Princeton development work was $61K; actual expenses for the work in Q3 and Q4 totaled $60K. Development effort at Fermilab was associated with the development and loading of the Runs DB; the Runs DB will be used to load SEGUE imaging data in early 2006. For the year, the revised forecast for SEGUE development work is $95K, or $43K (31%) below than the baseline budget.

Expenses associated with the DA upgrade were $39K against a baseline budget of $30K and included the efforts of personnel at Fermilab involved in the final elements of system commissioning and integration. For the year, the revised forecast for the final cost of the upgrade project is $256K, or $15K (6%) greater than the baseline budget of $241K.

A modest amount of effort went into the calibration effort in Q4. Expenses were approximately $5K, against a budget of $5K. The forecast for the year remains unchanged at $10K. The calibration effort is expected to ramp up over the course of the survey; consequently, the budget is larger in future years when efforts from available resources shift from development to calibration work.

Miscellaneous ARC corporate expenses (i.e., audit fees, bank fees, petty cash, and APO trailer rentals) were as expected. For the year, the revised forecast is $16K against the baseline budget of $32K.

7.3. Q4 Performance - Management Reserve

No management reserve funds were expended in Q4, or during the year. Unspent management reserve has been carried forward into the operating budget for future years.

7.4. Fiscal Year 2005 Cost Summary

The estimated value of in-kind contributions was $459K against anticipated contributions of $355K. The level of in-kind support provided by JHU, UW, UC and MSU/JINA was as agreed. The level of in-kind support provided by USNO was slightly less than anticipated. Fermilab contributed more than anticipated. The level of effort provided to support development, data processing and data distribution in 2005 was approximately 35% greater than planned. In addition, the salary costs of the individuals performing the work were greater than anticipated. We anticipate that the level of effort provided by Fermilab in future years will be closer to the baseline plan. Differences in salary costs will be taken into account when the operations budget forecast is revised.

Total cash expenses in 2005 were $1,979K against a baseline budget of $2,305K. No management reserve funds were expended in 2005. Unspent funds will be carried forward into the operating budget for future years, once final invoices are received from the institutions performing work in 2005.

8. PUBLICATIONS

In Q4, there were 20 papers based on SDSS data that were published by members of the SDSS collaboration. There were also 33 papers published by individuals outside of the collaboration, using publicly available data. Exhibit 3 lists papers published by members of the SDSS Collaboration; Exhibit 4 lists papers published by individuals outside of the SDSS collaboration.

At the time of this writing, there are 965 published refereed papers that include 'SDSS' or 'Sloan' in their title and/or abstract. These papers have been cited a total of 27,362 times, including 51 papers cited more than 100 times and 139 papers with 50 or more citations. In addition, there are 1336 un-refereed papers with "SDSS" or "Sloan" in the title and/or abstract.

Exhibit 1. CY2005 Cash Budget Forecast

SDSS-II CY2005 Budget Forecast as of January 31, 2006 *(in $000s)*

		Inst	Qtr 3 Jul-Sep Baseline Budget	Qtr 3 Jul-Sep Actual Expenses	Qtr 3 Jul-Sep Variance (%)	Qtr 4 Oct-Dec Baseline Budget	Qtr 4 Oct-Dec Actual Expenses	Qtr 4 Oct-Dec Variance (%)	CY2005 Total Baseline Budget	CY2005 Total Actual Expenses	CY2005 Total Variance (%)
OPERATIONS BUDGET - CASH EXPENSES											
1.0 Survey Management											
SSP-221	ARC Secretary/Treasurer	ARC	3	2	-23%	3	4	29%	6	6	3%
SSP-234	ARC Business Manager	ARC	16	14	-15%	16	19	19%	32	33	2%
SSP-246	PU Office of the Project Scientist	PU	62	57	-7%	5	1	-86%	66	58	-13%
SSP-248	FNAL Support for Survey Management	FNAL	14	8	-44%	14	14	-5%	29	22	-24%
SSP-267	UC Support for Survey Management	UC	45	29	-34%	14	12	-15%	59	42	-29%
SSP-291a	ARC Support for Public Affairs	ARC	0	0	---	0	2	---	0	2	---
SSP-291b	ARC Support for Spokesperson	ARC	0	0	---	0	0	---	0	0	---
SSP-291c	ARC Support for Collaboration Affairs	ARC	2	0	-100%	2	0	-100%	3	0	-100%
SSP-291i	ARC Support for Public Information Officer	ARC	7	1	-80%	7	4	-34%	13	6	-57%
	Survey Management Sub-total		**148**	**112**	**-24%**	**61**	**56**	-8%	**208**	**168**	**-20%**
2.0 Survey Operations											
2.1 Observing Systems											
SSP-231	UW Observing Systems Support	UW	68	54	-20%	68	40	-41%	136	95	-31%
SSP-232	PU Observing Systems Support	PU	2	13	471%	1	14	1005%	4	27	661%
SSP-242	FNAL Observing Systems Support	FNAL	71	87	24%	70	85	21%	141	172	22%
SSP-261	FNAL Observing Programs and DA Support	FNAL	5	0	-100%	3	4	17%	8	4	-50%
SSP-291d	ARC Observing Systems Support	ARC	44	11	-74%	44	47	7%	88	59	-33%
	Observing Systems Sub-total		189	166	-12%	187	190	2%	376	356	-5%
2.2 Observatory Support											
SSP-235	NMSU Site Support	NMSU	394	367	-7%	394	438	11%	787	805	2%
2.3 Data Processing											
SSP-240	FNAL Software and Data Processing Support	FNAL	76	56	-27%	76	78	3%	152	134	-12%
SSP-238	PU Software and Data Processing Support	PU	30	16	-47%	30	28	-7%	60	44	-27%
SSP-239	UC Software and Data Processing Support	UC	13	11	-9%	13	9	-25%	25	21	-17%
	Data Processing Sub-total		119	83	-30%	119	116	-3%	237	199	-16%
2.4 Data Distribution											
SSP-240	FNAL Software and Data Processing Support	FNAL	16	0	-100%	16	16	0%	31	16	-50%
SSP-268	FNAL Data Distribution Support	FNAL	8	0	-100%	8	12	46%	16	12	-27%
SSP-237	JHU Data Archive Development and Support	JHU	0	0	---	0	1	---	0	1	---
	Data Distribution Sub-total		24	0	-100%	24	28	20%	47	28	-40%
2.5 ARC Support for Survey Operations											
SSP91f	ARC Additional Scientific Support	ARC	3	0	-100%	3	0	-100%	5	0	-100%
SSP91h	ARC Observers' Research Support	ARC	3	0	-100%	3	0	-100%	6	0	-100%
	ARC Support Sub-total		6	0	-100%	6	0	-100%	11	0	-100%
	Survey Operations Sub-total		**730**	**616**	**-16%**	**728**	**771**	6%	**1,458**	**1,387**	**-5%**

Exhibit 1. CY2005 Cash Budget Forecast (continued)

SDSS-II CY2005 Budget Forecast as of January 31, 2006 *(in $000s)*

			Qtr 3 Jul-Sep			Qtr 4 Oct-Dec			CY2005 Total		
		Inst	Baseline Budget	Actual Expenses	Variance (%)	Baseline Budget	Actual Expenses	Variance (%)	Baseline Budget	Actual Expenses	Variance (%)
OPERATIONS BUDGET - CASH EXPENSES											
3.0 New Development											
3.1 SEGUE Survey Development											
SSP-138	PU Software and Data Processing Support	PU	61	21	-65%	0	33	---	61	55	-10%
SSP-137	JHU Data Archive Development	JHU	5	0	-100%	5	0	-100%	10	0	-100%
SSP-140	FNAL SDSS-II Development	FNAL	12	14	18%	0	0	---	12	14	18%
SSP-268	FNAL Data Distribution Support	FNAL	28	0	-100%	28	25	-9%	55	25	-55%
SSP-291J	ARC Support for SDSS-II Development	ARC	0	1	---	0	0	---	0	1	---
	SEGUE Development Sub-total		106	37	-65%	33	58	80%	138	95	-31%
3.2 Supernova Survey Development											
SSP-131	UW Supernova Survey Development	UW	15	8	-47%	0	0	---	15	8	-47%
SSP-139	UC Supernova Survey Development	UC	41	5	-88%	0	0	---	41	5	-88%
	APO Computer Room Upgrade	ARC	34	47	39%	0	0	---	34	47	39%
SSP-140	FNAL Supernova Survey Development	FNAL	72	14	-80%	0	0	---	72	14	-80%
	Supernova Development Sub-total		162	74	-54%	0	0	---	162	74	-54%
3.3 Data Acquisition System Upgrade											
SSP-161	FNAL Data Acquisition Upgrade	FNAL	211	218	3%	30	39	29%	241	256	6%
	DA Upgrade Sub-total		211	218	3%	30	39	29%	241	256	6%
3.4. Photometric Calibration Development											
SSP-138	PU Software and Data Processing Support	PU	5	5	0%	5	5	0%	10	10	0%
	Photometric Calibration Sub-total		5	5	0%	5	5	0%	10	10	0%
	New Development Sub-total		**484**	**334**	**-31%**	**68**	**102**	51%	**551**	**436**	**-21%**
4.0 ARC Corporate Support											
SSP91e	ARC Corporate Support	ARC	19	4	-80%	13	13	-6%	32	16	-50%
	ARC Corporate Support Sub-total		**19**	**4**	**-80%**	**13**	**13**	-6%	**32**	**16**	**-50%**
Cash Budget Sub-total			**1,381**	**1,065**	**-23%**	**869**	**914**	5%	**2,250**	**1,979**	**-12%**
5.0 Management Reserve		ARC	28	0	-100%	28	0	-100%	55	0	-100%
TOTAL CASH BUDGET			**1,408**	**1,065**	**-24%**	**897**	**914**	2%	**2,305**	1,979	**-14%**

Exhibit 2. CY2005 In-Kind Contribution Forecast

SDSS-II CY2005 Budget Forecast as of January 31, 2006 *(in $000s)*

			Qtr 3 Jul-Sep			Qtr 4 Oct-Dec			CY2005 Total		
		Inst	Baseline Budget	Actual Expenses	Variance (%)	Baseline Budget	Actual Expenses	Variance (%)	Baseline Budget	Actual Expenses	Variance (%)
OPERATIONS BUDGET: IN-KIND											
1.0 Survey Management											
SSP-248	FNAL Support for Survey Management	FNAL	33	44	32%	33	45	35%	66	89	34%
	Survey Management Sub-total		**33**	**44**	**32%**	**33**	**45**	**35%**	**66**	**89**	**34%**
2.0 Survey Operations											
2.1 Observing Systems											
SSP-231	UW Observing Systems Support	UW	16	16	3%	16	16	3%	31	32	3%
SSP-257	USNO Support for Survey Operations	USNO	4	0	-100%	4	0	-100%	8	0	-100%
	Observing Systems Sub-total		20	16	-18%	20	16	-18%	39	32	-18%
2.3 Data Processing											
SSP-239	UC Software and Data Processing Support	UC	0	0	---	0	6	---	0	6	---
SSP-240	FNAL Software and Data Processing Support	FNAL	2	53	---	2	92	---	4	144	---
SSP-257	USNO Support for Survey Operations	USNO	10	10	0%	10	0	-100%	19	10	-50%
	Data Processing Sub-total		12	62	439%	12	98	430%	23	160	597%
2.4 Data Distribution											
SSP-237	JHU Data Archive Development and Support	JHU	0	0	---	32	36	13%	32	36	13%
SSP-240	FNAL Software and Data Processing Support	FNAL	1	0	-100%	1	0	-100%	2	0	-100%
SSP-268	FNAL Data Distribution Support	FNAL	9	0	-100%	9	6	-34%	18	6	-67%
	Data Distribution Sub-total		10	0	-100%	42	42	-100%	52	42	-19%
	Survey Operations Sub-total		**41**	**78**	**90%**	**73**	**156**	114%	**114**	**234**	**106%**
3.0 New Development											
3.1 SEGUE Survey Development											
SSP-269	MSU SEGUE Software Development and Support	MSU	16	16	1%	16	16	1%	31	31	1%
	SEGUE Development Sub-total		16	16	1%	16	16	1%	31	31	1%
3.2 Supernova Survey Development											
SSP-139	UC Supernova Development and Support	UC	49	49	0%	0	0	---	49	49	0%
SSP-140	FNAL SDSS-II Development and Support	FNAL	95	56	-42%	0	0	---	95	56	-42%
	Supernova Development Sub-total		144	105	-27%	0	0	---	144	105	-27%
	New Development Sub-total		**160**	**120**	**-25%**	**16**	**16**	1%	**175**	**136**	**-22%**
TOTAL IN-KIND CONTRIBUTIONS			**234**	**242**	**4%**	**122**	**217**	78%	**355**	**459**	**29%**
TOTAL OPERATING BUDGET (Cash and In-kind)			**1,642**	**1,307**	**-20%**	**1,019**	**1,131**	11%	**2,660**	2,438	**-8%**

Exhibit 3. Papers from within the SDSS Collaboration

1. The Southern Flanking Fields of the 25 Orionis Group. AJ submitted - Peregrine M. McGehee, et al.

2. Constraining the Projected Radial Distribution of Galactic Satellites with the Sloan Digital Sky Survey. ApJ submitted - Jacqueline Chen, et al.

3. The RASS-SDSS Galaxy Cluster Survey. VII. On the Cluster Mass to Light Ratio and the Halo Occupation Distribution. A&A submitted - P. Popesso, et al

4. Constraining the Evolution of the Ionizing Background and the Epoch of Reionization with z~6 Quasars II: Analysis Using a Sample of 19 Quasars. AJ submitted - Xiaohui Fan, et al.

5. Galaxy Halo Masses and Satellite Fractions from Galaxy-Galaxy Lensing in the SDSS: Stellar Mass, Luminosity, Morphology, And Environment Dependencies. MNRAS submitted - Rachel Mandelbaum, et al.

6. SDSS J1534+1615AB: A Novel T Dwarf Binary Found with Laser Guide Star Adaptive Optics and Implications for the L/T Dwarf Transition. ApJ submitted - Michael C. Liu, et al.

7. The Milky Way Tomography with SDSS. ApJ submitted - Mario Juric, et al.

8. Gas Infall and Stochastic Star Formation in Galaxies in the Local Universe. MNRAS submitted - Guinevere Kauffmann, et al.

9. Cataclysmic Variables from SDSS V. The Fifth Year (2004). AJ accepted - Paula Szkody, et al.

10. SDSS J0806+2006 and SDSS J1353+1138: Two New Gravitationally Lensed Quasars from the Sloan Digital Sky Survey. AJ accepted -Naohisa Inada.

11. A Catalog of Spectroscopically Selected Close Binary Systems from the Sloan Digital Sky Survey Data Release Four. AJ submitted - Nicole M. Silvestri, et al.

12. A Survey of $z>5.7$ Quasars in the Sloan Digital Sky Survey IV: Discovery of Seven Additional Quasars. AJ accepted - Xiaohui Fan, et al.

13. Seventy New L and T Dwarfs from the Sloan Digital Sky Survey. AJ accepted - Kuenley Chiu, et al.

14. The SDSS Quasar Survey: Quasar Luminosity Function from Data Release Three. AJ submitted - Gordon Richards.

15. Very Small-Scale Clustering of Luminous Red Galaxies. ApJ submitted – Morad Masjedi, et al.

16. The Bright End of the Luminosity Function of Red Sequence Galaxies. MNRAS accepted - Yeong Loh, et al.

17. Cataclysmic Variables from SDSS IV. The Fourth Year (2003). AJ 129:2386 (2005) - Paula Szkody, et al.

18. The White Dwarf Luminosity Function from SDSS Imaging Data. AJ accepted - Hugh Harris, et al.

19. Spectral Variability of Quasars in the Sloan Digital Sky Survey. II: The CIV Line. ApJ accepted - Brian C. Wilhite, et al.

20. A search for the most massive galaxies: Double Trouble? AJ accepted - M. Bernardi, et al.

Exhibit 4. Publications Based on Public Data

1. The Luminosity-Weighted or 'Marked' Correlation Function. MNRAS submitted - Ramin Skibba, et al.

2. Galaxies in SDSS and DEEP2: a quiet life on the blue sequence? ApJ submitted - Michael Blanton, et al.

3. Morphology Segregation of Galaxies in the Color-Color Gradient Space. ApJL accepted - Changbom Park, et al.

4. The RASS-SDSS Galaxy CLuster Survey. V. The X-ray Underluminous Abell Clusters. A&A submitted - P. Popesso, et al.

5. What Galaxies Know About Their Nearest Cluster. ApJ submitted - A. Quintero, et al.

6. A Spectroscopic Study of the Ancient Milky Way: F- and G-Type Stars in the Third Data Release of the Sloan Digital Sky Survey. ApJ accepted - Allende Prieto, et al.

7. The Dependence of the Pairwise Velocity Dispersion on Galaxy Properties. MNRAS submitted - Cheng Li, et al.

8. The Dependence of Clustering on Galaxy Properties. MNRAS submitted - Cheng Li, et al.

9. Discovery of a Magnetic White Dwarf/Probable Brown Dwarf Short-Period Binary. ApJL submitted - G. Schmidt, et al.

10. The soft X-ray properties of quasars in the Sloan Digital Sky Survey. MNRAS submitted - Shiyin Shen, et al.

11. The RASS-SDSS Galaxy Cluster Survey. VI. The dependence of the cluster SFR on the cluster global properties. A&A submitted – P. Popesso, et al.

12. Acoustic oscillations in the SDSS Luminous Red Galaxy sample power spectrum. A&A submitted – Gert Huetsi (0507678).

13. Origin of [OII] Emission in Red Sequence and Post-starburst Galaxies. ApJ submitted - Renbin Yan, et al.

14. Lyman-alpha forest-CMB cross-correlation and the search for the ionized baryons at High Redshift. MNRAS submitted - Rupert A.C. Croft, et al.

15. Simulating Wide-Field Quasar Surveys from the Optical to Near-Infrared. MNRAS accepted – Natasha Maddox, et al.

16. An excess of damped Lyman alpha galaxies near QSOs. MNRAS accepted – David M. Russell, et al.

17. Objective Subclass Determination of Sloan Digital Sky Survey Unknown Spectral Objects. ApJ submitted – David Bazell, et al.

18. Selecting damped Lyman-alpha systems through CaII absorption. I: Dust depletions and reddening at z~1. MNRAS accepted – Vivienne Wild, et al.

19. The Discovery of Tidal Tails Around NGC 5466. ApJL submitted – V. Belokurov, et al.

20. Nature and completeness of galaxies detected in the Two Micron All Sky Survey: MNRAS submitted – Daniel H McIntosh, et al.

21. Optical Star-Formation Rate Indicators. ApJ, accepted – John Moustakas, et al.

22. Candidate Isolated Neutron Stars and Other Optically Blank X-ray Fields Identified from the ROSAT All-Sky and Sloan Digital Sky Surveys. AJ accepted – Marcel A. Aqueros, et al.

23. The chemical composition of metal-poor emission-line galaxies in the Data Release 3 of the Sloan Digital Sky Survey. A&A accepted – Y.I. Izotov, et al.

24. Semi-empirical analysis of Sloan Digital Sky Survey galaxies: II. The bimodality of the galaxy population revisited. MNRAS submitted – Abilio Mateus, et al

25. Effects of galaxy interactions in different environments. MNRAS accepted – M. Sol Alonso, et al.

26. Double Lobed Radio Quasars from the Sloan Digital Sky Survey. AJ accepted – W.H. de Vries, et al.

27. The Evolution of the Broad-Line Region among SDSS Quasars. A&S accepted -Tohru Nagao, et al.

28. Polar Outflows in Six Broad Absorption Line Quasars. ApJ submitted – Hongyan Zhou, et al.

29. Damped Lyman Alpha Systems at z<1.65: The Expanded SDSS HST Sample. ApJ accepted – Sandhya M. Rao, et al.

30. A Catalog of Edge-on Disk Galaxies: From Galaxies with a Bulge to Superthin Galaxies A&A accepted – S. J. Kasutsch, et al.

31. Properties of Galaxy Groups in the SDSS: I. The Dependence of Colour, Star Formation, and Morphology on Halo Mass. MNRAS submitted – S. M. Weinmann, et al.

32. Shapes of Clusters and Groups of Galaxies: Comparison of Model Predictions with Observations MNRAS submitted – D.J. Paz, et al.

33. A Census of Object Types and Redshift Estimates in the SDSS Photometric Catalog from a Trained Decision-Tree Classifier. AJ accepted – A.A. Suchkov, et al.