Sloan Digital Sky Survey

Quarterly Progress Report

Fourth Quarter 2000

February 6, 2001
Revised May 10, 2001

Table of Contents

1. OBSERVATION STATISTICS

1.1. Summary

The fourth quarter consisted of three dark runs nominally starting on October 18, 2000 and ending on January 4, 2001. The start of the first dark run was postponed until November 1 because of small delays in the scheduled realuminization of the primary mirror following the dedication on October 5, and one week of unscheduled maintenance to relubricate the rotator bearing. The latter required a disassembly and reassembly of the 2.5-m telescope and more details of this work are provided in Section 3.2. The late start reduced the duration of the first dark run to only six nights. Of these six nights, the first was spent completing shakedown tests to ensure that all systems were operating properly; the rest were lost to bad weather. The second dark run of the quarter (Nov/Dec) provided limited opportunities for imaging because of the weather. Since the observing systems were operational throughout this dark run, it was possible to obtain valuable spectroscopic data. The observing time in the third dark run (Dec/Jan) was reduced by both hardware unavailability and inclement weather. At the beginning of the third dark run (Dec/Jan), four nights were lost due to mirror control problems, and half of an additional night was dedicated to an extended shakedown. Finally, the imager was unavailable for observation for seven nights at the end of the dark run. Of these seven nights, only one night of potential imaging was lost. Inclement weather and the bright moon permitted only limited spectroscopic observations during six of the seven nights.

Another unfortunate consequence of the limited opportunity to obtain imaging data during Q4 is that the plate inventory appropriate for spectroscopy during December, January, and February sky remains small. At certain times during the night, therefore, the only sensible use of the observing systems is to repeat spectroscopic observations with plates that had previously been declared done. While such repeat exposures are important for testing and can also lead to interesting new science (e.g. studying variability in the spectra of quasars), this situation inhibits progress in covering the sky. It needs to be noted that no imaging had been done during December and January of the preceding two years and as a result the inventory of plates for the area of sky available in December and January was inadequate. Since very little new imaging data were obtained in Q4, the problem will continue to be with us for some time.

1.2. Time Use and Efficiency

Table 1.1 presents observing time in the same format as the same table in the Q3 report. The hours given for imaging and spectroscopy refer to the actual time exposing on the sky for science data. A new category, setup and calibration, which includes the total time expended for instrument change, plate changes, and set up of observations, was added to account for all scheduled time.

Table 1.1 Observing Time – Q4 2000 (units are in hours)

Dark Run	1 Nov – 6 Nov	11 Nov – 6 Dec	15 Dec – 4 Jan
Scheduled dark hours	68	230	234
Hours lost to weather	51	72	81
Engineering time	17	0	6
Equipment down time	0	11	48
Setup & calibration	0	65	40
Imaging	0	28	2
Spectroscopy	0	54	57

During some of the equipment down time, the weather was “cloudy”. In order to avoid counting the same hours twice, these hours show up only under equipment down time. Observing hours “lost to weather” means any weather-related reason that prevents observations, including clouds, blowing snow, and high winds. It does not include time when imaging is not possible because the sky is not photometric or the seeing is poorer than 1.5 arcseconds. Overall for the quarter, the amount of observing hours lost to weather was somewhat greater than average, although it is in line with expectations given the expected variability of the weather. After subtracting hours for inclement weather, engineering, and down time from the available hours, 56% of the time was used for science exposures during the second dark run and 60% of the possible time was used for science exposures during the third dark run. This is similar to the percentage of time used for science in Quarter 3.

A few other statistics are of interest. The nominal spectroscopic exposure time for good conditions is 45 minutes. If the seeing is poor, or if there are some cirrus clouds, the exposure time can be lengthened to meet the signal-to-noise ratio criterion. In Q4 the successful plates required an average of 60 minutes to complete. This increment of 15 minutes per plate gives a measure of the effect of these adverse conditions during Q4.

Not all of the plates exposed on the sky were declared done by the signal-to-noise criterion. A total of 55 plates were declared done (31 new plates and 24 repeated plates). There were, however, 111 attempts to expose plates, or a yield of 55 / 111 = 50%. We will investigate what factors led to this poor yield.

1.3. Status of Photometric Telescope Secondary Patches

Table 1.2 summarizes the status of the collection of the secondary patches in Q4. These are needed to carry out the photometric calibration of the imaging data.

Table 1.2. PT Secondary Patches Collected in Q4

	Done and verified	Done but not verified	Have old data	Total done	Not done, but high priority
2001 Jan 22	348	88	49	485	51

Explanation of categories:

Done, verified: Data have been obtained, processed, and verified to be of good quality.

Done, not verified: Data have been obtained but have not yet been processed.

Have old data: These are patches for which overlapping data were obtained before the layout of patches on the sky was revised in summer 2000. The old data are adequate for calibrations.

Total done: The sum of the above three categories.

Not done, high priority: This category includes patches not yet observed that either overlap existing data, patches that overlap anticipated new data, or patches for which repeat observations are needed for quality analysis purposes.

The statistics in Table 1.2 cover both the North and South hemispheres. The collection of patches in the South is complete, except for occasional repeats. In addition, for the first time "Not Done, High Priority" includes patches for areas that are high priority for future scanning with the 2.5m system. Considerable progress was made in reducing the number of patches in the "Not Done, High Priority" category from 120 at the beginning of the quarter to 51 at the end, in spite of the inclement weather. During the quarter, 254 patches were completed.

1.4. Summary of Imaging Observations to Date

Table 1.3 provides the status of imaging observations at the beginning and end of Q4. A similar set of statistics were presented in the Q3 report. Table 1.3 also corrects a typographical error in the table in the Q3 report. The categories in Table 1.3 have now been expanded to properly reflect progress toward the Southern Survey goals. The imaging part of the Southern Survey (roughly September through November) consists of three non-overlapping, 2.5-degree wide stripes. One is along the celestial equator, one is about 15 degrees north of the equator, and the third is about 10 degrees south of the equator. These have areas of 270, 205, and 270 square degrees, respectively. The latter two "outrigger" stripes are scanned in the same way as the stripes in the Northern Survey. The rationale for these stripes is that with a modest investment of imaging time, we achieve a substantial leverage in the ability to measure large-scale structure in the Southern hemisphere.

The central (equatorial) stripe, however, will be scanned as many times as practical during the course of the Survey. This strategy will provide a unique ability to discover objects that vary in intensity or in position over timescales from a day to 5 years. Moreover, once the scans are co-added, the depth of the photometric catalogues will be significantly deeper, enabling photometric redshifts to be determined for galaxies to even higher redshifts as well as numerous other extensions of the Northern Survey. The increased depth will not be uniform across the central stripe because some segments will have more repeat scans than others. The metric appropriate for judging the progress in the Southern central stripe is what we will call the repetition factor. This is computed by taking the total amount of imaging data acquired within the footprint for this stripe in the "Good" category and dividing that by the footprint area (270 square degrees).

It is important to know how the imaging survey is done in order to appreciate some of the complications (and nomenclature). Area of sky is covered by scanning along a pre-defined set of great circles. In the end, the imaging survey will cover a footprint of sky which is elliptical in outline and which avoids obscured directions in the Milky Way. A particular observation is called a "run": it is a part of one of the great circles, the beginning and end of which is determined by the date of observation, the specific time on that night the atmosphere permitted good seeing, etc. Runs will be overlapped to avoid leaving holes, and the amount of repeat scans needs to be accounted for. Eventually the whole of the great circle within the footprint is covered.

In more detail, the camera consists of 6 columns of detectors, each of which scans a narrow swath called a scanline. Each scanline is 2048 pixels wide; since a pixel is 0.4 arcsec, a scanline is 13.6 arc minutes (0.23 deg) wide. The scanlines are separated by gaps that are 90% of the size of a detector, such that a second run can fill in the gaps with a bit of margin to spare. (This margin is needed because the telescope tracking is done open-loop.) A set of six scanlines is called a strip; the interlaced pair of strips is called a stripe. Having both strips of a stripe is necessary for the subsequent selection of spectroscopic targets, which means that the area covered by the stripes at any point in the survey is the "coin of the realm."

Just as the scanlines overlap by about 10% of their width, so too the stripes overlap by a few arc minutes for the same reason: to allow for margin in the telescope pointing and tracking. The overlap between neighboring stripes varies depending on position along them because they are each centered on a great circle: the system of great circles of course converges at the poles. The minimum overlap is at the equator of the system of great circles; we defined the system such that stripes are separated by exactly 2.5 degrees there.

In summary, while the camera actually observes 20.5 square degrees per hour in terms of pixels on the sky, the effective rate of covering new sky is 9% less, or 18.75 square degrees per hour. The details of stripe-to-stripe overlap depend on the final geometry for the survey footprint, i.e. how close to the poles each stripe gets.

Not all of the imaging data obtained meet the quality standards of the survey - some data are rejected because the image quality is sub-standard, or because there is evidence that the sky was not uniformly transparent, or for some other reason. Thus as the data move through the pipelines, some data are rejected and that part of the sky will need to be re-observed.

To capture the nature of the way the data are collected and processed, we have devised a hierarchy of categories, each of which is tracked. These categories, which are used in Table 1.3, are defined as follows:

1. Gross – Total area scanned, as reported in idReport files. This is the sum of the area of the science runs that are good enough to be candidates for data processing. In the context of SDSS data processing, this should be the area that one computes by scanning all the idReport files in the "golden directory."

2. Net – "Total" area that has been successfully processed through the image pipelines. The excluded area includes ramp-up frames and the frames were rejected due to bad tracking, clouds, setup time, or bad seeing. Net area is the area of sky that is stored in the Operations Database.

3. Good – "Net" area that passes QA tests for seeing, tracking, and photometricity. This statistic is not reduced by overlaps.

4. Unique – "Good" area that is corrected for overlaps between separate runs on the same stripe and that lie within the official survey area. This is one of the statistics that is used to track progress against the baseline and it was the statistic that was used to define the baseline of January 2000.

5. Footprint – "Unique" area after removing overlap of stripes. This is ultimately the area of sky that the survey has covered and it is the statistic for the revised baseline that most realistically measures progress toward the survey goals for the North Galactic Cap and the southern outrigger stripes.

6. The Repetition Factor - The ratio of the sum of the Good areas scanned in a given footprint to the area of the footprint. It is one of the two statistics that are needed to measure progress toward the survey goals for the Southern Equatorial stripe. The other is the footprint statistic.

Table 1.3. Summary of Imaging Data to Date

	Gross	Net	Good	Unique	Footprint	Repetition Factor
North Galactic Cap						
through 2000 Oct 01	1545	1366	1150	978	838	
through 2001 Jan 10	1825	1574	1294	1098	935	
Southern Outrigger Stripes						
through 2000 Oct 01	616	577	484	374	374	
through 2001 Jan 10	775	719	586	422	422	
Southern Equatorial Stripes						
through 2000 Oct 01	406	377	135	51	41	0.41
through 2001 Jan 10	569	536	284	201	169	0.96

1.5. Summary of Spectroscopic Observations to Date

The status of spectroscopic observations is measured by the number of plates that meet the signal to noise (s/n) requirements and other quality measures. A plate normally provides 640 spectra. The cumulative statistics for spectroscopic observations are presented in Table 1.4.

Table 1.4. Summary of Spectroscopic Observations to Date

	Complete Plates	Repeat Plates
North Galactic Cap		
through 2000 Oct 01	58	1
through 2001 Jan 10	80	4
Southern Outrigger Stripes		
through 2000 Oct 01	6	0
through 2001 Jan 10	18	4
Southern Equatorial Stripe		
through 2000 Oct 01	36	0
through 2001 Jan 10	40	15

Through the end of Q4 (i.e., for CY2000), we obtained 80 plates in the North and 58 plates in the South. These numbers count those plates that 1) meet a uniform S/N criterion; 2) are within the survey footprint; and 3) are science plates as opposed to test plates. Of these, in Q4 alone we obtained 22 plates in the North and 16 plates in the South.

1.6. Comparisons to the Baseline Projections

The original baseline projection was given in an Appendix to the Q1 2000 Report. This is revised in draft form as an Appendix to this Report with particular elaborations related to the southern stripes. Also, instead of using 20.5 square degrees per hour to correct from imaging time to square degrees, we realize now that it is better to use 18.75 square degrees per hour as described in Section 1.4. The new baseline reflects this change.

1.6.1. Imaging Statistics

In Q4, we obtained 120 square degrees of new "unique" data in the North, 48 square degrees of new "unique" data on the South outrigger stripes, and 150 square degrees on the south equatorial stripe. In all cases, the amount of new data collected was below the baseline, as shown in the comparison presented in Table 1.5. The values listed under Baseline come from the Appendix; the values listed under Actual are from the "Unique" column of Table 1.3.

Table 1.5. Summary of Imaging Data Obtained in 2000-Q4 and Cumulative-to-Date

	Baseline			Actual (Unique)		
	North	South (outriggers)	South (equatorial)	North	South (outriggers)	South (equatorial)
2000-Q4 only	384	242	240	120	48	150
Total through 2001 Jan 10	563	336	335	1098	422	201

With regard to imaging progress to date, we are ahead of the imaging baseline in all areas except for the southern equatorial stripe. Figure 1.1 graphically shows the imaging progress in the Northern Survey Region. The main conclusion is that despite the fact that we are officially ahead or close to par, the recent rate of obtaining new imaging data is not adequate (cf. the slope of the plot).

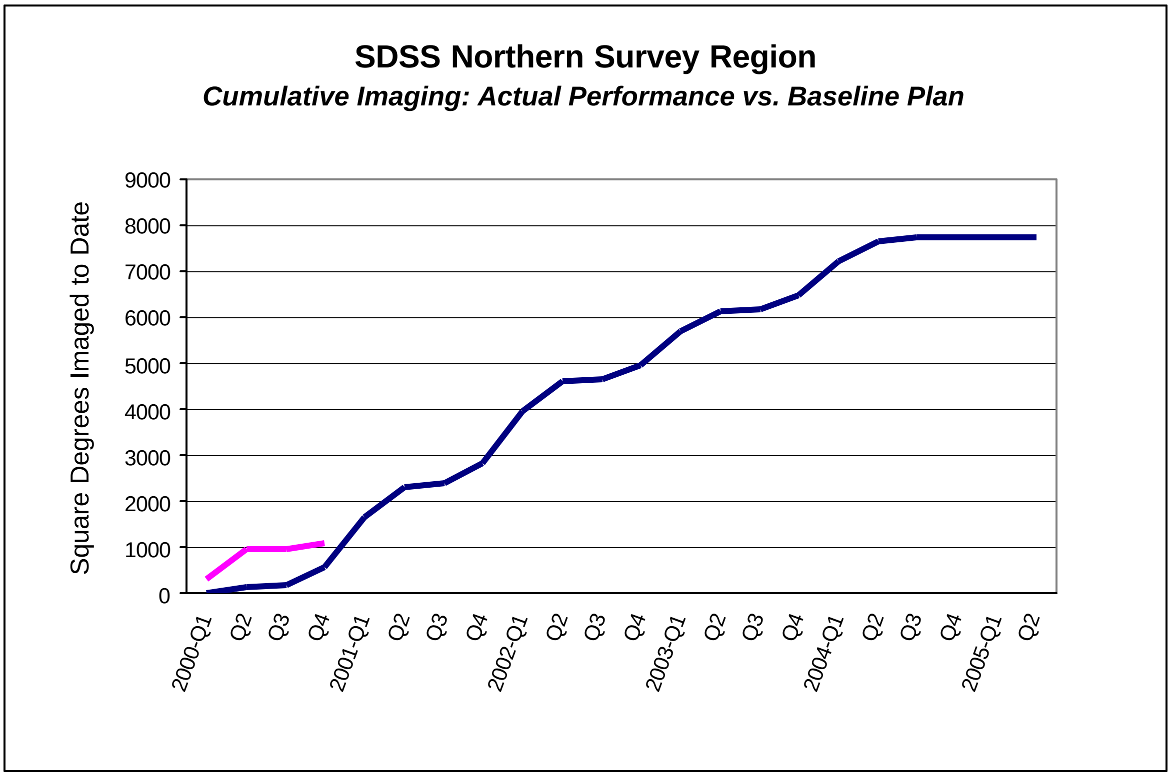

Figure 1.1. Cumulative Imaging Data vs. Baseline Plan

1.6.2. Spectroscopy Statistics

The baseline projection called for zero plates in the North (because of an assumed phase lag between imaging and spectroscopy) and 61 plates in the South at the end of December. As just mentioned, the actual numbers were 80 and 58, respectively. The spectroscopic phase of the survey is on schedule in spite of inclement weather. Figure 1.2 graphically shows spectroscopic progress in the Northern Survey Region. Again, while we are ahead of the baseline projection, the rate at which we are observing plates is not adequate.

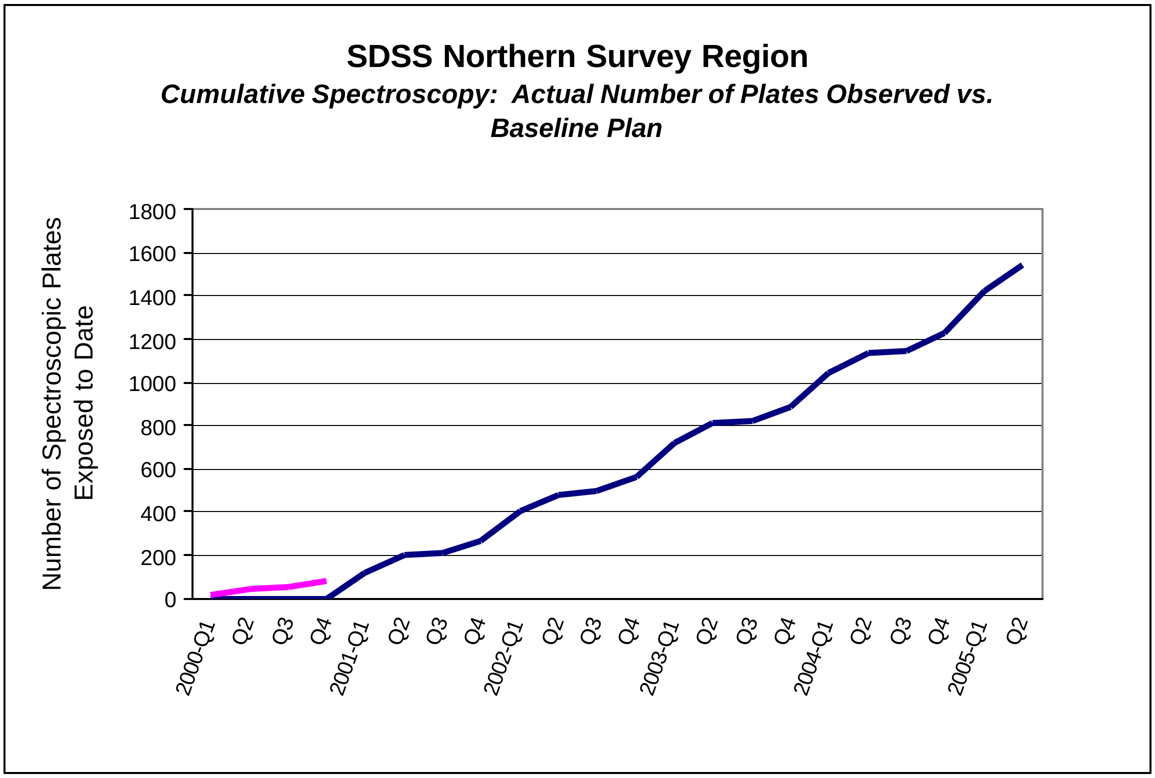

Figure 1.2. Cumulative Spectroscopy vs. Baseline Plan

1.6.3. Conclusions

The purpose of the baseline projection is to help us evaluate our progress. After 12 months of operations, we can conclude the following:

1. Weather seems to be roughly in line with expectations, although the fraction of the time that the weather was photometric with good seeing as opposed to the fraction of time that the weather permitted spectroscopy has been disappointing. We note that the seeing must be better than 1.5 arcseconds before imaging data meet requirements, whereas as long as the seeing is better than 3 arcseconds it is suitable for spectroscopy.

2. For CY2000, the fraction of the time lost to engineering and down time was about 12%, only slightly worse than the expectation of 10% that is in the baseline. However, this did not include the longer-than-planned scheduled down time in October.

3. The overhead in spectroscopic observations (cartridge exchange, calibrations, field set-up) is in line with the baseline assumptions (see the Q3 Report for an accounting), at least for the best nights. We can conclude that the baseline assumptions are not unrealistic, but we need to improve operations so that low overhead becomes the norm, not the exception.

4. The fraction of time that could have been used for exposure on the sky that actually was so used, mentioned in Section 1.2, compares favorably with the baseline. There are two reasons for this: the baseline had a generous allowance for inefficiency at the start of the survey (ramping up even in Q1 of 2001), and not all of the data obtained in the time intervals shown in Table 1.1 actually resulted in good data. Good data were actually obtained in only 56% of the spectroscopic time and 60% of the imaging time. In the future, we will fail to meet the baseline efficiency goals unless we improve our efficiency.

5. The ratio of time spent undertaking imaging to time spent undertaking spectroscopy has evolved from 2.4 in Q1 to 0.3 in Q4, essentially reflecting the maturing of the spectroscopic systems in 2000 and the existence of imaging data useful for spectroscopic targeting. It also reflects the sky conditions, which often prevented imaging. The ratio of imaging time to spectroscopic time evolved qualitatively but not in detail according to the baseline, but there are no issues related to this difference other than the ones already mentioned.

6. We are spending on average 60 minutes exposure per successful plate, whereas the baseline assumed 45 minutes per plate. We have been able to partly offset the increased exposure time by using some moonlit time for spectroscopy (which the baseline did not assume). Nonetheless, we recognize that we must work to reduce the average exposure time per plate in order to meet the baseline spectroscopy goals.

7. The thermal environment around the telescope is affecting image quality, which consequently affects the amount of time that the telescope is capable of meeting performance requirements. Improving the thermal environment around the telescope will increase the fraction of time that image quality will be better than 1.2 arcseconds. It will also result in shorter spectroscopic exposures.

1.7. Target Selection and Plate Drilling

We did not report statistics that are appropriate for target selection and plate drilling last quarter. However, in the third quarter we drilled 60 plates covering two "chunks" of sky, all in the south. In the fourth quarter, we drilled a total of 69 plates in two drilling runs. These plates cover four "chunks" of sky, three in the north and one in the south. All plates passed QA testing and were shipped to APO for observing.

2. DATA PROCESSING AND DISTRIBUTION

2.1. Imaging

The imaging pipelines run in the following order: ssc, astrom, photo, nfcalib. The following paragraphs summarize fourth quarter progress and performance.

<u>ssc</u>: This pipeline extracts images of bright stars for the run to be used for the astrometric calibration and to characterize the point spread function. During Q4, v4_6 was released, which included the following maintenance fixes. First, duplicate entries in the output files that occurred

from two closely spaced stars were removed. Second, we trap and gracefully handle an error condition that occurs when wrapping around RA=360.

astrom: This pipeline determines position calibration for the run. During Q4, v3_3 was released, which implemented ad hoc solutions for g' residuals versus color, thus improving our quality control sensitivity.

photo: This pipeline detects and measures objects. During Q4, releases v5_2_9 through v5_2_14 implemented fixes for robust factory operations.

mtPipe: This pipeline produces calibration fields from the monitor telescope data. No updates were necessary for this pipeline during the quarter.

nfcalib: This pipeline performs the flux calibration of objects. During Q4, releases v1_7 through v1_7_7 implemented a factory-ready version of this pipeline.

After the data is processed through nfcalib, target selection is performed. The most recent version of Target is v2_13_4. In November, the first "survey ready" version was declared, with final algorithms and parameters set for galaxy and QSO selection. In December, the galaxy magnitude limit was changed from 17.67 to 17.77, per a decision of the Change Control Board.

2.2. Spectroscopy

The spectroscopic pipelines run in the following order: idlspec2D, spectro. The following paragraphs summarize fourth quarter progress and performance.

idlspec2d: This pipeline extracts, calibrates, and combines the red and blue halves of all of the spectra in a set of exposures. A production version, v3c, was declared in May, and releases v4_1_2 through v4_4_0 in Q4 added spectrophotometry (flux calibration), correction of near infrared scattering, better handling of telluric corrections, and some bookkeeping improvements.

spectro: This pipeline identifies and measures the calibrated spectra. A production version v4 was declared in July, and releases v5_0_2 through v5_2_1 in Q4 improved error handling, made a more robust determination of the redshift in cases where multiple estimates are possible, and improved the galaxy/QSO identification and line fitting.

2.3. Data Distribution

It was difficult to keep the science database loaded with current versions of reductions for several reasons. Problems with the commercial software corrupted the database on several occasions. Additionally, we now have multiple reductions of some runs. We will handle this by having two instances of the science database. One will contain the version of the catalogs used to drill plates, and the other will contain the most recent version of the catalogs.

2.4. The Data Processing Factory

Significant progress was made towards a "hands-off" operation of the pipelines. Much of the work, from spooling raw imaging data through loading the operational database, now occurs with the automation software. We have also automated the target selection and plate design process. Another significant change was to move processing of spectro data to a farm of PCs running Linux, necessitated by a dramatic increase in the CPU and memory requirements of the programs. This transition is essentially complete and in the validation process. The task list to build the factory is 75% complete.

In the fourth quarter, we processed imaging data in a timely fashion for the plate drilling runs that began on November 6 and December 12. The total amount of data processed in the quarter was reported above in Section 1.

3. OBSERVING SYSTEMS

The fourth quarter was spent gathering a relatively small amount of survey data (owing to inclement weather and the events described below) and pursuing a very aggressive (and not always anticipated) engineering program, mostly centered around the telescope but involving the camera vacuum system as well late in the quarter.

3.1. Controlling the Thermal Environment around the Telescope

Early in the quarter, a general plan for mitigating the very unfavorable thermal situation was evolved, and we began to pursue it aggressively. The problem, in summary, has three components:

First, we are dissipating far more heat in the lower enclosure than was originally anticipated, which causes the temperature in the lower enclosure to be considerably warmer than ambient. This has three bad effects. The instrument fluid-cooling chiller, which is located on the rotating rack in the lower enclosure, cannot cope with the high temperature in the lower enclosure. As a result, it delivers 'coolant' to the instruments that is considerably hotter than the outside temperature, which in turn causes the instruments to run hot. Additionally, the steel support structure for the telescope terminates in the lower enclosure and so the steel structure runs hot. Lastly, the air from the lower enclosure is exhausted to the east (usually downwind) and, because it is hot, can and does cause serious seeing disturbances.

Second, the upper enclosure air handling system is inadequate to prevent serious temperature stratification in the daytime, and cannot cool the enclosure to match the very low opening temperatures experienced in the winter.

Third, the system that pulls air through the mirror and into the lower enclosure is not working nearly as it should. This is mostly due to leaks in the lower enclosure allowing outside air to come in, but is also due to unanticipated restrictions in the airflow associated with wiring and hoses in the telescope. The result is that the time constant for the mirror to equilibrate with the ambient air after opening is very long, and if the nighttime temperature changes rapidly, the equilibration never occurs.

During the fourth quarter, we made progress on all three aspects. We began a program of switching from air-cooling to liquid cooling (with the energy dumped into an existing ground-loop system) for equipment for which this is possible in the lower enclosure. We have begun implementing a system of large switching DC power supplies to DC-DC converters for equipment both downstairs and on the telescope, and, most importantly, we are replacing the very power-hungry pumps that supply compressed air and vacuum to the mirror support system with a new system of pumps that will be located in the operations building. These changes will be completed in the first quarter of 2001, and should result in lowering the power dissipated in the lower enclosure and the resulting temperature rise by about a factor of three.

The stratification will be dealt with by better ducting in the upper enclosure, but we have not reached a consensus on exactly what to do here. It appears from very recent experiments that we can get away with something quite simple and inexpensive, and this will be implemented as soon as we are sure. The low-temperature performance (or rather, lack thereof) will be much more expensive to deal with, and we will not implement something in time to deal with this winter, but will decide on a system, do the engineering, and get it installed before the onset of cold weather next year.

The mirror ventilation system is being dealt with by eliminating leaks in the lower enclosure--this is essentially finished, but in the process we have become painfully aware of just how difficult it is to make a building airtight--and adding high-pressure fans in the telescope support cone to help pull air directly through the mirror ventilation tubes. In addition, a number of air leaks in the telescope have been discovered and fixed.

The situation is currently much better than it was at the end of the summer, but we still have a long way to go. We expect a large improvement when the power elimination program is finished in the first quarter of 2001, and another when we get the fans in, but there are some issues associated with vibration and acoustic problems here, and it will be necessary to proceed cautiously. By mid-spring we expect to see the full benefit of this program, and expect to be rid of our low-temperature problems by the next time we might encounter them.

In addition to attempting to relieve the problems, we have installed an extensive system of solid-state thermometers in the mirrors and on the telescope structure. The current state of this system is that 100 of 108 thermometers are installed and connected, and the system almost works; there are systematics associated with the control/digitizing circuitry that we are currently trying to understand and fix. We have learned from the system so far that our worst fears about the mirror relaxation time are realized and that the image quality degradation we see is thermally induced. We strongly suspected these things, but it is very useful and encouraging that the large thermal program we have undertaken will certainly bear fruit in the end.

3.2. The Telescope

The primary mirror was removed from the telescope and aluminized in early October, immediately following the dedication. We got a superb coat. While the mirror was out, previous serious problems with high torque in the rotator bearing were investigated and it was discovered that the upper seal in this large preloaded bearing had failed and that the bearing had almost certainly become contaminated as a result. We were lucky that we were able to get field service from Rotek, the European manufacturer of the bearing. The rotator bearing was removed, disassembled, inspected, cleaned, reassembled, and reinstalled with new seals. In the process, we learned that the bearing was in fact severely contaminated but undamaged. We also discovered that one of the rotator drive motor bearings was showing signs of failure, so we replaced this bearing as well. The rotator work was accomplished expeditiously and uneventfully but took us out of service for about a week longer than we had anticipated for the aluminization shutdown.

Work proceeded on various 'clean-up' tasks ranging from finishing the telescope-windscreen bumpers to a fairly major rework of the interlock system in anticipation of finally getting the drive roller slip-detection monitoring in place and getting a somewhat de-scoped instrument-change protection system in place. The development of the latter has required a great deal of planning effort during the quarter. We anticipate that the slip-detection system will be installed in March and the instrument-change system will be finished by June.

3.3. The Instruments

The fourth quarter saw the installation of the autofill system for the spectrographs, which keeps the 10-liter intermediate liquid-nitrogen dewars full during the day and when not observing (they last through a night). After some difficulty with the circuitry, they appear to be working well.

The camera was disassembled in the summer of 2000 to fix various small problems and worked well through Q3 and the first two dark runs of the fourth quarter, though there was clearly a problem with a small intermittent leak in the vacuum system. One of the things installed in the summer was instrumentation to measure and record the vacuum, so we do not know whether the phenomenon was new or not. Regardless, the leak in the vacuum system became bad enough in late December that we decided to take the camera out of service and repair the leak. After a scramble to assemble people and resources, the camera was disassembled and the leak (in a bad O-ring) repaired in early January. As earlier noted, we may have lost one night of imaging due to the camera repair.

3.4. The Observing Software

In general, the observing software programs performed well during the fourth quarter. Incremental upgrades were made to almost all systems. The results of this activity were that SOP, MOP, and the MCP have disposed of all their critical problem reports and many of the planned enhancements (some of this work was performed early in the present quarter). Work continues on IOP, which needs some major rework, and the TPM, which continues to respond to hardware upgrades. Work also continues on developing and upgrading the mountain tactical databases, which is going well. We have seen this quarter the regular operation of a piece of code called Son-of-Spectro, which is a 'lightweight' spectroscopic reduction pipeline for quick-look QA of spectroscopic data, and Hoggpt, a similar program for almost-real-time reduction of PT data to assess data quality and photometricity. The real-time QA tools in astroline for the camera data are adequate, but some improvements to enhance efficiency and the ability to judge data quality on the fly are yet to be done.

3.5. The Photometric Filters

We have had a problem with the filters that define our photometric bands that we have known about for some months, on which we have made substantial progress this quarter. The thin interference films which define the long-wavelength edges of the g', r', and i' filters were laid down by an evaporative process in vacuum by their manufacturer, Asahi, in Tokyo. It has become evident in the past few months, following preliminary measurement of the response shapes using the spectrophotometric instrumentation provided to us by the JPG, that these filters as installed in the camera (in which the films are in a hard vacuum) are considerably different from their state as delivered. After some investigation, it turns out that the films are slightly porous and hygroscopic, so that water from even very dry air is adsorbed and changes the refractive index of the films slightly. This water is not attached to the filters in the camera vacuum, and so the measurements of the filters in air at delivery are different from the shapes as installed in vacuum. The shifts are not large enough to cause significant trouble with the photometric system, but there are calibration issues, since the filters in the PT that are used for the calibration are kept in dry air, and the shift is not nearly so large as is seen in the camera. In the fourth quarter, we carefully measured the responses of all of the camera CCDs and have reassured ourselves that the 'shift-to-vacuum' is stable and consistent and that we can define an accurate, consistent photometric system for the survey data, albeit a slightly different one from that defined in the early published literature on the subject. The fear that the filters in the PT may well not be stable as the very small amount of water in the air purge for these filters changes and the temperature changes has led us to the decision to remake these using a newly-developed ion-beam deposition technique which eliminates the voids in the film. The new filters will be made to match as closely as possible the filters as they currently are in the PT in order to introduce the smallest systematics. We expect to have the new filters in hand this spring.

3.6. Photometry

It was recognized some months ago that photometric accuracy was going to be one of our major problems, and we have adopted a quite vigorous program to try to find and fix the problems. One aspect of this program was the establishment of an independent 'tiger team' to write an independent piece of reduction code to check the MT pipeline and to investigate whether the photometric calibration strategy was appropriate and satisfactory.

The result of the involvement of this team and accompanying intense effort on the part of many others is that the situation is much improved and, within the margin of error of determination of these things, probably barely meets our specs most of the time. Work on this problem will continue fairly intensely in the coming year, hopefully with results that are quite unambiguously satisfactory at the end. The improvements thus far mostly revolve around fixing previously inadequate techniques used in final calibration, improved techniques for rejecting marginal calibration data, and the installation of the Hoggpt photometric robot discussed above, which allows fast assessment of photometric quality by providing almost real-time photometric solutions.

Future work will involve investigations toward placing more reliance on the 2.5-meter photometric data for several important aspects of relative self-calibration and direct comparison with and correction of the PT calibration patches. On the basis both of experience and calculations, the 2.5-meter should be essentially as good an absolute photometric telescope as the PT is, and the ensemble of data should be used in as statistically powerful a manner as possible. Work will also be done on the use of the 2.5-meter in a large-pixel (binned) fast-scan mode to create a calibration network of its own across the survey area. In a few nights of clear weather but not good seeing, a completely independent calibration system can be established and tied to the standard star network through the whole ensemble of PT patches. There are several software and hardware issues associated with this 'Apache Wheel' approach, but these will be investigated and, we hope, disposed of in the coming year. At present, a bit of prototype data has been obtained and awaits processing. Another effort will be to quantify the assessment of the 10-micron cloud camera images to establish a priori the quality of a night, and to apply more sophisticated QA to PT data using the systematics of stellar colors. This program overall has already been quite successful, and we anticipate that a completely satisfactory system will exist within the year. The new calibrations will be applied retroactively to already-released data in order to maintain the homogeneity of the survey.

4. SURVEY PLANNING

Several incremental improvements were made to the tools used to plan the survey and track its progress.

1. A plug plate can only be observed for a limited time during the night due to effects of differential refraction. Up to now these limits were estimated only coarsely. The time limits are now calculated accurately, and for most plates, these limits have increased. This eases the job of scheduling spectroscopic observing during the night, increasing the efficiency somewhat.

2. The patch database, used to track and plan observations with the Photometric Telescope, was improved extensively so that it is now easier to track progress. A table listing the areas of sky that have been scanned has been added, ensuring that patches needed to calibrate imaging data are obtained in a timely fashion.

3. Programs to track imaging progress have been improved incrementally. It is now possible to generate web pages with succinct information on all scans that have been obtained, their quality, and their targeting status. The draft a monthly imaging plan has been improved so that it guarantees that short scans still cover calibration patches.

5. COST REPORT

The approved SDSS project budget for 2000 consisted of two parts: $1,848K for Fermilab, US Naval Observatory (USNO), and Los Alamos National Laboratory (LANL) expenses, which are paid by these institutions and counted as in kind contributions; and $3,700K for ARC funded expenses. Actual ARC funded expenses in 2000 were $3,592K. The ARC funded budget is summarized in Table 5.1, which compares fourth quarter and annual expenses against the baseline. A more complete table comparing actual to baseline performance is included as an attachment to this report.

Table 5.1. ARC-Funded 4th Quarter Expenses and Summary for 2000 ($K)

	2000 – 4th Quarter		2000 - Total	
Category	Baseline Budget	Actual Expenses	Baseline Budget	Actual Expenses
Project Management and Survey Coordination	11	49	244	301
Observing Systems	231	225	1,074	1,141
Data Processing and Distribution	156	169	686	689
Observatory Support	296	347	1,185	1,198
ARC – Corporate Expenses	49	57	139	242
Capital Improvements	0	0	0	21
Sub-total	743	848	3,327	3,592
Undistributed Contingency	71	0	372	107
Total	814	848	3,700	3,700

As noted in previous reports, final fourth quarter expenses were not available for all of the institutional budgets at the time this report was prepared, due to variations in the timeliness of the accounting systems for the various institutions performing work for the SDSS. In these instances, fourth quarter expenses have been estimated to the best of our ability to provide a forecast of total ARC funded expenses for 2000. The reported expenses will be revised as final invoices are received from the supported institutions.

5.1. Revised First, Second, and Third Quarter Expenses

Actual expenses for the first, second, and third quarters have been adjusted to reflect revised costs for these periods. Some of the adjustments involved moving costs from one quarter to the next to reflect when expenses were actually incurred against the various accounts. Other adjustments were made to replace forecasts made in previous quarters. In summary, these adjustments resulted in a first quarter decrease of $6K, a second quarter increase of $14K, and a third quarter decrease of $19K, for a net change of ($11K) compared to the expenses reported in the third quarter report.

5.2. Fourth Quarter Performance

The sum of in-kind contributions for the fourth quarter was $566K and was provided by Fermilab, Los Alamos, and USNO. Fermilab provided support for the data acquisition system at APO and data processing systems at Fermilab as agreed. Fermilab also provided the agreed upon level of support for survey management and for data processing operations, and a slightly higher than agreed upon level of effort to support the observing systems. This additional effort included work to mitigate thermal problems affecting telescope performance. Los Alamos provided additional observing support for Photometric Telescope operations, which represents an increase in the level of commitment that occurred after the baseline budget was established. USNO provided the agreed upon level of effort during the quarter.

The sum of ARC-funded expenses for the fourth quarter was $848K. This was $105K above the fourth-quarter budget, without contingency, of $743K.

Project management costs exceeded the Q4 budget for a number of reasons. First, summer salary expenses at Princeton were charged against the wrong account in the third quarter and this error was corrected in Q4. Since there were no summer salaries in the Q4 budget, the cost transfer to correct the error makes the Q4 budget look overspent. In fact, the budget for summer salaries was underspent for the year. Second, work on the spectroscopic commissioning effort at the University of Pittsburgh continued beyond that in the baseline plan. In accordance with the baseline plan, the budget to support the Spectroscopic Scientist at Pittsburgh was scheduled to end with the second quarter. Although the Spectroscopic Scientist met his performance goals in the first and second quarters, a significant portion of his budget was unspent. A proposal by the Spectroscopic Scientist to continue the work related to spectroscopic commissioning was approved by the Director and the unspent funds from the first half of the year were used to fund this work through the remainder of 2000. Thus, Q4 expenses exceed the budget simply because there was no Q4 budget for Pittsburgh. Notwithstanding, total Pittsburgh expenses for the year were within $500 of their allocated budget. Finally, the agreement established between Fermilab and ARC for 2000 stated that Fermilab would be reimbursed for up to $70K in travel expenses associated with the SDSS. In the baseline plan, the total anticipated cost of Fermilab travel exceeded $70K and so Fermilab anticipated providing a fraction of the fourth quarter travel costs as an in-kind contribution. However, since Fermilab travel expenses were only $65K for the year, $5K in survey management travel costs was charged against the ARC funded budget in the fourth quarter.

Data processing and distribution costs exceeded the Q4 budget for the reasons forecast in the third quarter report. The Princeton software development budget was overspent due to a higher than anticipated increase in salary rates and the Princeton cost for supporting the independent effort on the photometric calibration system. The latter represents a task that was added in July and thus not included in the baseline plan. Additionally, a new ARC-funded position was added at Fermilab to support observing operations and provide software development and maintenance support for the observing programs. It should be noted that the creation of this new position at Fermilab and the increase in the Princeton scope of work to support photometric calibrations were approved by the Change Control Board since these constituted increases over the baseline plan and budget.

Finally, the ARC Corporate budget exceeded the Q4 budget for the reasons noted and forecast in earlier reports. The first quarter report described a number of items that were not in the budget approved in November 1999. In the fourth quarter, these items included APO petty cash and cleaning expenses and AAS meeting expenses (the cost of a booth at the AAS meeting in San Diego and SDSS PR brochures). Excluding these additional expenses, fourth quarter ARC Corporate expenses were in line with the baseline plan.

5.3. Year 2000 Budget Performance Summary

The approved ARC funded budget for 2000 was $3,700K, including $372K for contingency. Actual expenses for the year were $3,592, which required using $265K of the contingency. Expenses that were paid for with contingency are described in Table 5.2. The remaining $107K of contingency not spent in 2000 will be carried forward.

Table 5.2. Increases in the 2000 Budget over the Baseline, by Category

Item	Cost Increase ($K)
ARC corporate expenses	115
Telescope and instrument support	67
Project management support	57
APO engineering support building design	21
Observatory support	13
Data processing and distribution support	3
APO dedication	(4)
ARC personnel	(7)
Total	265
Contingency adjustment	(265)
Undistributed contingency remaining	107

6. PUBLICATIONS

Statistical Properties of Bright Galaxies in the SDSS Photometric System
AJ submitted – K. Shimasaku, et al.

Galaxy Number Counts from the Sloan Digital Sky Survey Commissioning Data
AJ submitted – Naoki Yasuda, et al.

Weak Lensing Measurements of 42 SDSS/RASS Galaxy Clusters
AJ submitted – Erin Scott Sheldon, et al.

Detecting Clusters of Galaxies in the Sloan Digital Sky Survey I: Monte Carlo Comparison of Cluster Detection Algorithms
AJ submitted – Rita S.J. Kim, et al.

A New Very Cool White Dwarf Discovered by the Sloan Digital Sky Survey
ApJ Lett accepted – Hugh Harris, et al.

Colors of 2625 Quasars a $0<z<5$ Measured in the Sloan Digital Sky Survey Photometric System
AJ accepted – Gordon Richards, et al.

The Luminosity Function of Galaxies in SDSS Commissioning Data
AJ submitted – Michael Blanton, et al.

High-Redshift Quasars Found in Sloan Digital Sky Survey Commissioning Data V. Hobby-Eberly Telescope Observations.
AJ accepted – Don Schneider, et al.

Stellar Population Studies with the SDSS
ApJ Lett accepted – B. Chen, et al.

Optical and Infrared Colors of Stars Observed by 2MASS and SDSS
AJ, 120, 2615-2626, 2000 – Finlator Kristian, et al.

Revision Log

Revision	Date	Description
1	5/10/01	The text for Section 4 was added. It was not included in the version that was distributed on February 6, 2001.