Sloan Digital Sky Survey

Quarterly Progress Report

Third Quarter 2000

October 31, 2000

Table of Contents

1. OBSERVATION STATISTICS

In the third quarter, we imaged all three stripes of the southern Galactic cap. We also devised an explicit formula relating to the size of stellar images for determining when the imaging data can support target selection. For spectroscopy, the third quarter featured a refined implementation of the S/N criterion that ensured that all of the plates considered on the mountain to be properly exposed did, in fact, pass the quality criterion once the data were fully reduced. The new S/N criterion also allows us to expose during intervals of moonlight and even twilight. Finally, our best efficiency achieved this quarter matched the average efficiency necessary to reach the baseline goals.

1.1. Time Use and Efficiency

Observing time in the third quarter is less than other quarters by one dark run to accommodate maintenance during the monsoon season. Moreover, the SDSS dedication on October 5, and the disassembly of the telescope for re-aluminization of the primary mirror soon thereafter, shortened the number of available nights in this particular quarter by a few.

Table 1.1 gives an accounting of telescope usage in a format similar to previous quarterly reports. Note that we explicitly report the down time, or time lost due to equipment problems, and it is no longer relevant to report commissioning observations.

Table 1.1. Observing Time - Q3 2000 (units are hours)

Dark Run	22 Aug - 6 Sep	21 Sep - 3 Oct
Dark hours in run	112.0	110.0
Hours lost to clouds	47.5	21.3
Engineering time	3.0	0
Down time	6.0	8.4
Imaging	9.1	26.7
Spectroscopy	24.2	31.5

A measure of our observing efficiency can be derived by dividing the total number of hours spent for imaging and spectroscopy by the hours available in principle (dark hours minus cloud hours). Doing so gives 52% for Aug/Sept and 66% for Sept/Oct. These values are in line with the best of previous dark runs, although the accounting is not completely consistent since beginning with this quarter, some of the spectroscopy was undertaken in non-dark hours.

1.2. Baseline Projection

What really counts, of course, is how well we are doing with respect to the baseline. The baseline assumes that the observing systems will be working properly 90% of the dark, clear time. According to Table 1.1, adding the engineering time to the down time, we experienced 92% uptime. The baseline moreover assumes that 60% to 70% of the time the weather will permit at least spectroscopic observations (specifically, for Q3). In Q3, we experienced conditions that were judged useful 69% of the time, which is in line with what was anticipated.

The baseline projection calls for a ramp-up between 2000 and 2001, anticipating that we will become more efficient at collecting data. The 2000-Q3 baseline calls for 23 hours of imaging (one hour of imaging gives 20.5 square degrees, uncorrected for overlaps) and 16 plates, and the 2001-Q3 baseline calls for 42 hours of imaging and 67 plates. In fact, we obtained 36 hours of imaging and 53 plates in the third quarter, which exceeded the baseline and suggests that we will be able to meet the 2001-Q3 baseline.

1.3. Spectroscopy Details

Since time spent on spectroscopy is roughly 80% of observing time (in the northern Galactic Cap), and since the efficiency of spectroscopic operations is the most problematic, it is worth looking more closely at the time spent for spectroscopy (Table 1.2).

Table 1.2. Spectroscopic Time Use in 2000-Q3

Dark Run	22 Aug - 6 Sep	21 Sep - 3 Oct
Total hours spectroscopy	43.7	58.4
Number of plates	23	30
Exposure hours	24.2	31.5
Overhead hours	19.5	26.9

The first line defines the dates of the observing run. The second line accounts for elapsed time when spectroscopic operations were being undertaken. On average, the total time integrating on the sky was 1:03 (h:m) for each science exposure; multiplying this by the number of plates obtained gives the number of exposure hours. The overhead hours shown in Table 1.2 are the difference between the total hours and the exposure hours. It can be seen that the average time to expose a plate (including all overheads) is 1:55, and the average overhead time (i.e., not exposing on the sky) per plate is 0:52. This time is the sum of: calibration exposures, CCD read-out, diamond or smear exposures, cartridge exchange, and field acquisition.

The best spectroscopic efficiency was on September 29, when 4 plates were successfully exposed. All of these plates reached the requisite S/N with 0:45 exposure time, and the average overhead time was 0:37, for a total per plate of 1:22. In fact, this best-yet performance essentially equals the baseline assumptions for exposure time and overhead, a noteworthy achievement. Another achievement is that on October 3 we obtained 7 good plates in one night.

1.4. Hardware Downtime

For this quarter we have made an explicit accounting of time lost on the 2.5-meter telescope because of various problems. These are detailed in the following table.

Table 1.3. Time Lost due to Hardware Problems

MJD	Hours Lost	Problem
51780	1	Spectrograph sp2 dewar warming
51781	2.75	Poor telescope pointing - fiducials
51790	2.25	Interlock module failure
51810	3.5	M2 Galil communications failure
51812	1.5	M2 Galil communications failure
51818	2.0	Rotator aborts
Various	1.0	Rotator aborts on various nights
Total	14.0	

Each of these problems received prompt attention. Several were resolved in the third quarter and the outstanding problem with the rotator aborts will be addressed early in the fourth quarter.

In addition to the downtime reported in Table 1.3, the PT was down for four nights due to vacuum problems in the CCD dewar. The system has since been modified to preclude this problem in the future. The four lost nights affect three nights of imaging and we are investigating the implications for the photometric calibration of those data.

Some time is effectively lost because of thermal problems that affect seeing. One way to estimate this is as follows (albeit very crudely): as mentioned above, the average total integration time per plate was 63 minutes. 45 minutes is nominal, which means we exposed an extra 18 minutes per plate to get the requisite S/N. We can estimate that 2/3 of this extra time was allocated because of cirrus, and the remainder (6 minutes) for seeing. 53 plates x 6 minutes per plate = 5.3 hours. Of course, not all of this lost time due to seeing is within our control.

1.5. Rate of Observing Secondary Patches

For the first time we report here the status of collecting photometry of the secondary patches, which is of course necessary to support the calibration of the imaging scans. There are 1278 secondary patches in the Northern survey region, which require 688 hours of scanning time with the 2.5-m telescope. In other words, if for each hour of imaging time of the 2.5-m telescope we acquired 1.9 secondary patches with the PT, then we will have kept up. Since the PT can observe in poor seeing, this goal is even easier to achieve.

In the first dark run of Q3 we observed 81 secondary patches with the PT, and in the second dark run we observed 77 secondary patches. Elsewhere in this report we discuss the strategic issues associated with the secondary patches: namely, while the rate of obtaining patches is adequate, we have not necessarily acquired the right patches at the right time.

1.6. Summary of Imaging Observations to Date

Figure 1.1 plots the imaging data collected to date that has been corrected for overlapping scans and that lies within the official survey area against the baseline plan for the imaging survey. This data is defined as "unique" in Table 1.4, which presents a history of all of the imaging data collected to date.

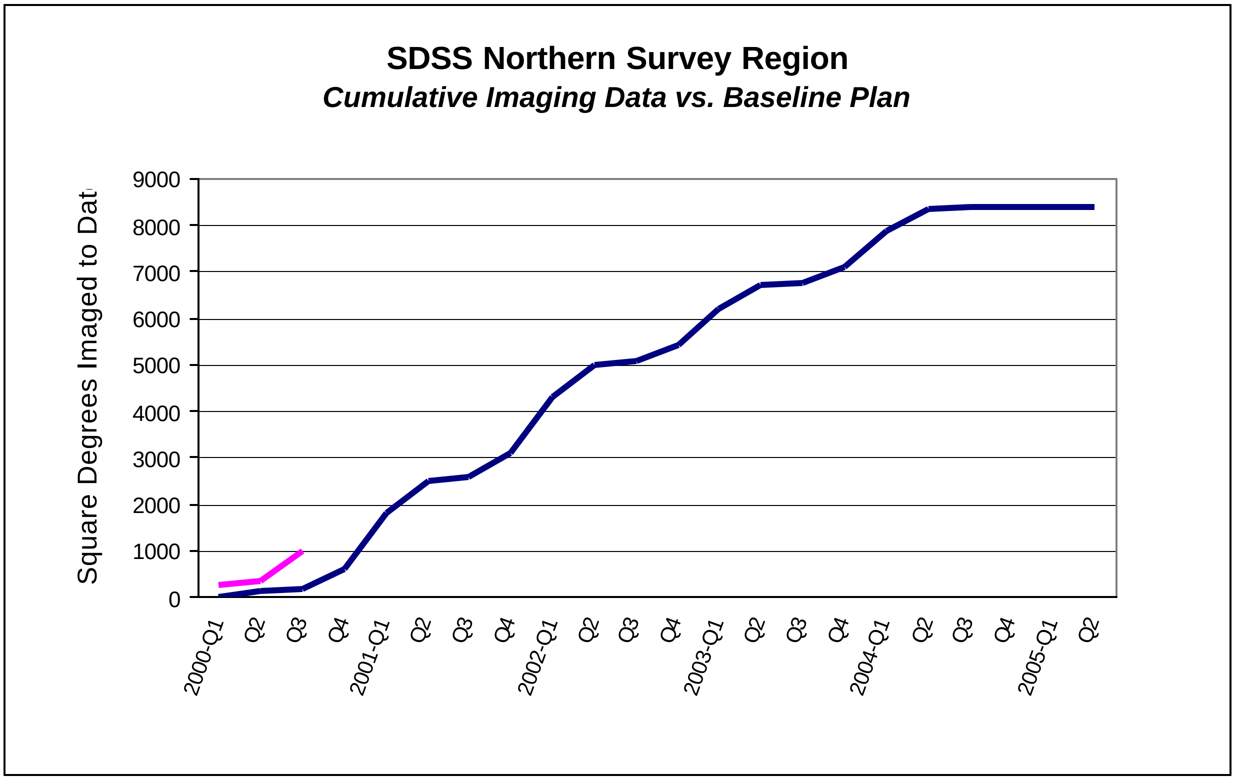

Figure 1.1. Cumulative Imaging Data vs. Baseline Plan

It should be noted that the numbers reported in Table 1.4 supersede previous accounts because we have now applied the recently revised seeing quality criterion. It should also be noted that in the South, the baseline plan calls for repeated scans of the equatorial stripe. Thus once the three southern stripes have been completed, the number of repeated scans of the equatorial stripe is the correct metric, not the cumulative area.

Table 1.4. Summary of Imaging Data to Date

	Gross	Net	Good	Unique	Footprint
NORTH					
1998 Jul 01	0	0	0	0	0
1998 Oct 01	0	0	0	0	0
1999 Jan 01	0	0	0	0	0
1999 Apr 01	422	394	394	260	224
1999 Jul 01	422	394	394	260	224
1999 Oct 01	422	394	394	260	224
2000 Jan 01	422	394	394	260	224
2000 Apr 01	631	581	510	364	315
2000 Jul 01	1717	1490	1231	992	860
2000 Oct 01	1490	1231	992	860	1717
SOUTH					
1998 Jul 01	0	0	0	0	0
1998 Oct 01	212	208	40	25	25
1999 Jan 01	212	208	40	25	25
1999 Apr 01	212	208	40	25	25
1999 Jul 01	212	208	40	25	25
1999 Oct 01	212	208	40	25	25
2000 Jan 01	541	503	209	165	165
2000 Apr 01	541	503	209	165	165
2000 Jul 01	541	503	209	165	165
2000 Oct 10	927	876	554	415	415

The table categories are defined as follows:

1. Gross - Total area scanned as reported in idReport files. These are the science runs that are good enough to run through data processing. This should be the area that one computes by scanning all the idReport files in the "golden directory".

2. Net - total area successfully processed through data processing. The excluded area includes ramp-up frames and whatever frames were rejected due to screwy tracking, setup time, or rotten seeing. This is the area that is actually put into opdb.

3. Good - net area that passes QA tests for seeing and astrometry.

4. Unique - good area that is corrected for overlapping scans and that lies within the official survey area. This is the number to use to track progress against some milestone. It is also the area that we use for planning future observations.

5. Footprint - Unique scans after removing overlap of stripes. This is ultimately the area of sky that the survey has covered.

1.6. Summary of Spectroscopic Observations to Date

In the first quarter report we said that of all the observed plates that passed the S/N criterion, 15 had been obtained in 2000-Q1 and an additional 6 had been obtained in 1999. In Q2 we reported 30 good plates. In section 1.3 above we reported 53 plates this quarter. Not only is there a clear trend to obtaining more good plates each quarter, but the yield (total number found to be good versus total number attempted) has also improved steadily.

2. DATA PROCESSING AND DISTRIBUTION

2.1. Imaging

Efforts this past quarter have focused primarily on reprocessing existing imaging data in order to design plates for the fall observing runs and to obtain a uniform set of data processed through a single version of all pipelines. The photometric pipeline was upgraded to a new version that fixes problems with photometry in poor seeing. Six imaging runs were processed and used to design about 66 plates. Several additional runs were processed; however, a problem was subsequently discovered in the new version of the photometric pipeline serious enough to warrant reprocessing these runs after the problem was fixed.

To augment the computing power available to process data, the data processing pipelines were ported to run on the Fermilab Linux "farms" - a cluster of workstations.

A major goal is to achieve better control of pipeline versioning. To this end, three testbeds of data are being set up, one of imaging data to test the imaging pipelines, one of object catalogs to test target selection, and one of spectroscopic plate data to test the spectroscopic pipelines. These testbeds will be used to compare the outputs of new pipeline versions with existing ones being used in production in order to validate them. A second goal is to automate as much as possible the running of pipelines, in essence producing a "factory." The steps that are automated include submission of jobs, monitoring the status of the jobs, and verification of outputs to determine if the next stage of processing can begin. In order to accomplish this work, the plate design run in October will be skipped.

2.2. Spectroscopy

Spectroscopic data continue to be processed immediately after they are acquired on the mountain. In addition, this quarter we re-processed data obtained in previous quarters with new versions of the Spectroscopic Pipeline.

2.3. Photometric Telescope

To date, we have observed and processed 230 secondary patches that are acceptable. About 20 of these patches will be re-observed at slightly different positions on the sky to improve the overlap with the 2.5-m scans. We are short about 120 patches if we wish to fully calibrate all the 2.5-m data in hand.

Minor bugs in the pipeline to process the patches were fixed and several cosmetic changes made to make the pipeline easier to run. Code was added to remove "fringing" patterns from the i' and z' bands; this code is not yet used in production. The layout of secondary patches on the sky was revamped to improve the overlap between the patches and the 2.5 m scans. Finally, work continued on evaluating the performance of the final calibration pipeline and improving its robustness.

2.4. Data Distribution

No new data were released to the collaboration while we focused on getting the testbeds set up and data reprocessed through a consistent set of pipelines.

Testing has begun on writing data to an Enstore tape robot. This robot has much greater capacity than the previous HPSS based system. However, the write speeds turn out to be too slow for the current archiving techniques. The Enstore developers will attempt to develop a strategy to increase the speeds.

Work has begun to make the SX database into a standard product that can be compiled and operated by Fermilab. This work remains in progress. A second database (the "chunk" database) was created to help deal with situations where a run has been processed multiple times. The "chunk" database contains the processing outputs used to select targets and design plug plates.

3. OBSERVING SYSTEMS

3.1. Thermal/Image Quality Issues

During the third quarter, we focused heavily on finding the causes of the less-than-perfectly-satisfactory image quality we have been obtaining in the imaging data thus far. Probable causes of the degradation fall broadly into three classes of problems: unsatisfactory seeing at the site, problems with the optics or their supports, and thermal problems that degrade either the optics or the seeing or both.

To address the first, several experiments were done with the DIMM, a small telescope that measures differential image motion in a way that can be correlated with seeing. The DIMM testing strongly indicates that the site seeing is significantly better than the image quality we typically obtain. In the near future, we plan to mount the DIMM permanently near the Sloan Telescope, which will not only allow us to monitor the improvements we make, but will be a powerful tactical tool that will help us make the decision whether to image or do spectroscopy.

The testing data that were delivered with the optics indicate that the optics are quite good, but of course the secondary was damaged last year and we have always been somewhat suspicious of the common corrector. We therefore contracted with the Lick Observatory optical shop to have both of these pieces re-measured this summer with their superb profilometer. The measuring was accomplished quite quickly but the test results were only just recently completed. The tests show that both pieces are very good; the corrector as good as was claimed originally and the secondary quite superb; no figure degradation apparently occurred as a result of the break.

While all of this was going on, we were gathering thermal data that began to point strongly to the trouble being due to thermal problems. We uncovered several problems that fall into two broad classes. The first is that we are generating too much heat on the telescope and especially in the lower enclosure. As a result, heat is conducted through the steel of the telescope and creates disturbances at the observing level, and the instruments run about 7C warmer than the ambient, which also creates serious disturbance in the air and likely causes trouble with the figure of the primary mirror. The second is that the primary mirror itself is a lightweight borosilicate glass structure whose optical quality depends on remaining isothermal and exchanging heat efficiently with the ambient air. Unfortunately, the ventilation system designed to accomplish this does not work satisfactorily for a variety of reasons. Fortunately, none of these problems are complex and most are not difficult and we have already begun addressing them in a phased approach.

3.2. Equipment Protection

Substantial progress has been made in addressing several equipment protection issues remaining from last spring. The issue of safely handling the spectroscopic corrector lens has been fully addressed; a housing and special-purpose cart was built and installed and works very well. The bumpers that prevent collisions between the telescope and windbaffle were constructed and fully installed on the elevation axis. Lightning-protection hardware was reworked, with all conductors from outside now isolated with optical fiber except for power, which enters the building with considerable surge and flashover protection. A metal shroud was installed over the enclosure power rail system to mitigate lightning concerns. Remaining lightning protection issues include full implementation of an early warning system and secure grounding of the movable building. Other equipment protection issues include final implementation of the azimuth collision bumpers and the slip-detection system that protects the drive disks from capstan slippage. The azimuth bumpers are scheduled for completion in the fourth quarter and the slip detection system scheduled for implementation in the first quarter of 2001.

3.3. Pointing and Tracking

One of the most pressing efficiency issues has been the quality of telescope pointing; the 2.5-meter must point with an RMS accuracy of about 2-3 arcseconds to allow fast acquisition of spectroscopic fields and satisfactory alignment of imaging scans. To address this, we have satisfactorily implemented a system that corrects the pointing at fifteen-degree intervals.

Outstanding issues include tracking and frictional problems with the instrument rotator. With regard to tracking, the rotator does not track correctly when asked to move quickly and there is a strong suspicion that the error is proportional to the tracking rate and is always present. We believe this is a software issue in that the rotator goes where it is told, but it is not told correctly. We are investigating this and anticipate resolving the problem during the fourth quarter. We also have a new issue with the rotator that arose in the September-October dark run, namely that it has developed a severe frictional problem when moving at very low speeds. The torque required to move it slowly rises until the servo overloads and aborts. A series of tests will be run in the very near future to isolate the problem and point, we hope, to a solution.

3.4. Photometry

In the third quarter, we established a "tiger team" to investigate the issue of photometric calibration, which we saw as a major impediment for the successful full-up operation of the survey. Results from the coordinated effort of this group with the Fermilab photometric group indicate that the site, photometric telescope, and software will suffice to calibrate the survey within the scientific requirements. Another set of problems we face is a recently discovered instability in the filters we are using. The interference films are apparently slightly hygroscopic and change their refractive index substantially when going from ordinary room air to vacuum. Once in vacuum, they appear to be very stable, but the cutoff wavelengths vary substantially with temperature in air. The camera filters are in vacuum and all of the calibration filters are in air and so this whole area is the subject of concentrated research at present. There is no real evidence that these effects have caused serious accuracy problems, but it is almost certain that the "system" we are using is affected by them and will have to be corrected.

4. SURVEY PLANNING

During the third quarter, criteria were developed and approved by the Change Control Board to accept or reject data based on seeing. The requirement is essentially that the full width at half maximum (FWHM) of a star not exceed 1.5 arcsecs averaged over the entire 30 CCDs of the array.

4.1. Observing Aids

Several new observing tools have been provided to APO to aid in validating data.

1. A seeing analysis tool aids in determining if imaging observations are acceptable. A new program analyzes the positions of stars in near-real time and determines if the telescope pointing and rotator orientation are correct. This decreases the setup time for imaging by several minutes per scan.

2. Son of Spectro was written and installed; this program analyzes spectroscopic exposures in near real time and determines if they have adequate signal/noise. This gives the observers the necessary information to determine the optimum amount of time to expose a plate.

3. The orientations of the guide fibers in 8 of 9 cartridges were determined and incorporated into the guiding program. This information makes the setup for spectroscopic observing much faster.

A significant effort is underway to improve the tools needed for planning the survey, tracking observations as they are taken, and providing feedback from data processing to planning. Two databases are now used at APO to track observations as they occur.

1. The plate inventory database has been in place for a year and works well at tracking which plates are in stock, which have been observed, and, once data processing has occurred, which can be retired or need to be re-observed.

2. A second database was installed this quarter and is used by the PT to track its observations of patches. There are limited mechanisms in place to flag observations as being either good or bad. This database also has the capability to store information about 2.5-m scans and flag the patches needed to calibrate those scans as being high priority for observing next.

Three programs are in various stages of development to aid in planning observations.

1. The stripe planning program accepts as input the current area of sky imaged and outputs a monthly plan for imaging that optimizes the time to completion of the imaging portion of the survey. This tool exists but is in need of much improvement.

2. The plate layout program accepts as input the current list of tiles and plates that have been designed, a list of plates that have been observed, and a list of new tiles output by target selection, and determines the exact parameters for drilling new plates. Because of refraction effects, a plate must be designed for a specific time of night. This program makes sure that an adequate supply of plates is available throughout a night for one or two dark runs in the immediate future.

3. The plate planning program accepts all the inputs of the previous two programs and determines which runs of imaging data already in hand are optimally placed for the next round of plate design and drilling. This program aids in determining which imaging runs should have high priority for data processing.

4.2. Target Selection

The major activity for target selection was verification and enhancement of the QSO target selection code. Spectra of quasar candidates from 40 plates were examined by eye and used to determine the efficiency of the selection algorithm. The code was modified extensively, mainly to reject objects not likely to be real quasars. The testing of this code has been limited by the lack of data that were homogeneously processed and calibrated. The target selection testbed described in Section 2 is being set up to alleviate this problem.

Improvements to galaxy and quasar codes are largely in the direction of rejecting unlikely candidates from the pool of targets. This means that plates designed using the new code will include virtually all good targets that would have been selected with the old code, with interlopers being removed.

4.3. Photometric Calibrations

The goal is to achieve 2% rms accuracy in the g', r', and i' filters, and 3% rms in u' and z'. This accuracy is being achieved on occasion but not always. One of the biggest problems at present is having a proper supply of good secondary patches; the shortage is due primarily to an improper layout of patches on the sky last year. Fortunately, the PT can observe at a rate that should allow us to catch up.

5. COST REPORT

The approved SDSS project budget for 2000 consists of two parts: $1,848K for Fermilab, the 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. The current forecast of ARC funded expenses for 2000 is $3,700K, with $31K of undistributed contingency remaining. This is an improvement over the second quarter report, wherein we forecast an $18K overrun in the budget.

The ARC funded budget is summarized in Table 5.1 and presented in detail in Appendix 1. Both tables compare third quarter expenses and the current forecast against the baseline.

Table 5.1. ARC-Funded 3rd Quarter Expenses and Revised Forecast ($K)

	2000 – 3rd Quarter		2000 - Total	
Category	Baseline	Actual Expenses	Baseline	Forecast
Project Management and Science Direction	63	100	244	312
Observing Systems	224	236	1,074	1,152
Data Processing and Distribution	166	182	686	738
Observatory Support	296	268	1,185	1,185
ARC – Corporate Expenses	37	76	139	259
Capital Improvements	0	21	0	21
Sub-total	787	883	3,327	3,668
Undistributed Contingency	71	0	372	31
Total	858	883	3,700	3,700

5.1. Revised First and Second Quarter Expenses

Actual expenses for Observatory Support in the first and second quarters were revised to reflect final costs and the actual split of expenses between these quarters. These adjustments resulted in a first quarter increase of $46.9K and a second quarter decrease of $38.7K. The net result is that total expenses through the second quarter are $8.2K higher than that reported in the second quarter report.

5.2. Third Quarter Performance

Fermilab, LANL, and the USNO provided the agreed upon level of effort throughout the quarter and both plan to provide the same level of effort for the remainder of 2000. Fermilab also provided support for the data acquisition system at APO and data processing systems at Fermilab as agreed.

The sum of ARC-funded expenses for the third quarter was $883K. This is $96K above the third-quarter baseline, without contingency, of $787K. Third quarter expenses were higher than the baseline for the following reasons. First, project management costs exceeded the baseline because 1) summer salaries had been incorrectly distributed evenly throughout the year in the baseline budget when they should have been concentrated in the summer; and 2) work on the spectroscopic commissioning effort at the University of Pittsburgh that was not planned for in the baseline was funded using funds left over from the underspent Pittsburgh budget from the first and second quarters. Although these items resulted in a cost increase against the third quarter baseline, the total budget for the year remains unchanged. Second, the Observing Systems budget was slightly overspent due largely to the procurement of spare components and replacement parts that were not budgeted for in the baseline and the need to redesign the slip detection system as a result of tests made on the prototype system. Third, the Data Processing and Distribution budget was overspent due to an increase in salary rates and the addition of costs to support the independent effort on the photometric calibration system. Fourth, the ARC Corporate budget exceeded the baseline for several of the reasons noted and forecast in the second quarter report. In the third quarter, costs not in the baseline included the profilometry measurements made by Lick Observatory; procurement of a shroud to mitigate lightning protection concerns on the telescope enclosure power feed system; and procurement of new workstations for use by the observers at APO. In addition, the design work was completed on the proposed APO Engineering Support Building; this design work has provided us with a solid cost estimate and set of bid documents that we can use when funds become available for this capital improvement. Finally, there is a cost savings to report. The Observatory Support budget was underspent by $28K due in part to the resignation of two observers and in part to the way in which vacation time is budgeted and charged by NMSU. This cost savings was projected in the second quarter and helped to offset the cost overrun in the Observatory Support budget in the second quarter.

5.3. Year 2000 Budget Forecast

The current forecast for the year is $341K above the baseline, without contingency. This is less than the forecast overrun of $391K presented in the second quarter report and can be covered by applying a significant portion of the $372K in contingency contained in the approved budget. Table 5.2 outlines the distribution of projected costs contributing to the forecast increase.

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

Item	Projected Cost ($K)
ARC corporate expenses	120
Telescope and instrument support	78
Project management support	68
Data processing and distribution support	52
APO engineering support building design	21
APO dedication	8
ARC personnel	(6)
Subtotal	341
Contingency/reserve adjustment	(341)
Undistributed contingency remaining	31

6. PUBLICATIONS

The First Hour of Extra-galactic Data of the Sloan Digital Sky Survey Spectroscopic Commissioning: The Coma Cluster
AJ submitted - Francisco Castander

Detection of Massive Tidal Tails around the Globular Cluster Pal 5 with SDSS Commissioning Data
AJ accepted - Michael Odenkirchen

High-Redshift Quasars Found in Sloan Digital Sky Survey Commissioning Data IV: Luminosity Function from the Fall Equatorial Stripe Sample
AJ accepted (Jan 2001) - Xiaohui Fan

High-Redshift Quasars Found in Sloan Digital Sky Survey Commissioning Data III: A Color Selected Sample at $i<20$ in the Fall Equatorial Stripe
AJ accepted (Jan 2001) - Xiaohui Fan