SDSS-III Project Execution Plan

Version 1.1

February 5, 2009

This is version 1.1 of the Project Execution Plan for SDSS-III, submitted to the National Science Foundation on February 6, 2009. This document draws on broad input from members of the SDSS-III survey teams. The principal authors and reviewers of the document text, who are responsible for its content, are:

Daniel Eisenstein, University of Arizona, Director

David Weinberg, Ohio State University, Project Scientist

Bruce Gillespie, Johns Hopkins University, Program Manager

David Schlegel, Lawrence Berkeley National Laboratory, BOSS PI

Steven Majewski, University of Virginia, APOGEE PI

Connie Rockosi, Lick Observatory, SEGUE-2 PI

Jian Ge, University of Florida, MARVELS PI

Jim Gunn, Princeton University, Observing Systems Lead and Infrastructure Lead

Mark Klaene, Apache Point Observatory, Site Operations Manager

Donald Schneider, Pennyslvania State University, Survey Coordinator

Michael Blanton, New York University, Data Coordinator

Jordan Raddick, Johns Hopkins University, Education and Public Outreach Coordinator

Michael Evans, University of Washington, ARC Business Manager

Revision History

Version 1.0: Submitted to NSF on December 31, 2008

Version 1.1: Updated Chapter 2 with text on APOGEE-MARVELS co-observing, more precise definitions of metrics. Added Appendix A on forecasts against these metrics. Added WBS, schedule, and budget as additional appendices, contained in separate files from this Word document. Added data distribution milestones section. Assorted cosmetic changes.

1. Organizational Structure

This section describes the management of the third phase of the Sloan Digital Sky Survey (SDSS-III) under the auspices of the Astrophysical Research Consortium (ARC). It describes the roles and responsibilities of ARC, the ARC Board of Governors, the Advisory Council, the Director, the Project Scientist, the Program Manager, and other senior personnel serving in key positions within the project.

The flow of accountability from the Board through the Advisory Council to the Director and the rest of the top survey management is shown in Figure 1.1.

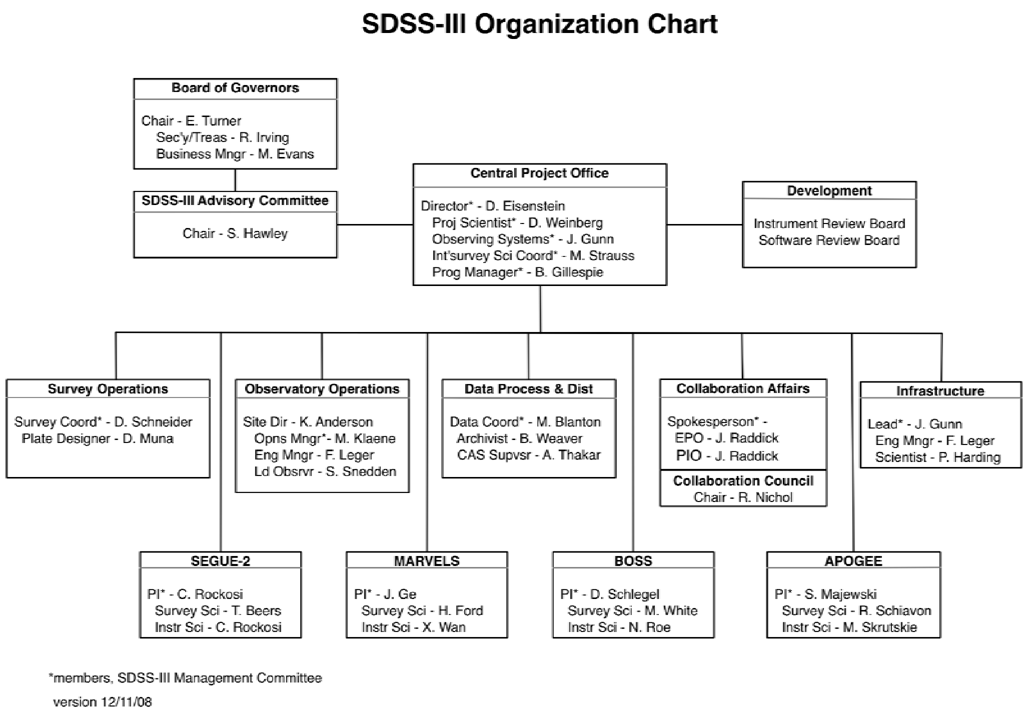

Figure 1.1. Organization Chart for ARC / SDSS-III Management

1.1. The Astrophysical Research Consortium

The Astrophysical Research Consortium (ARC) was created to provide the faculty, staff, and students from its member institutions access to modern astronomical equipment for their research and educational programs. ARC owns and operates the Apache Point Observatory (APO) near Sunspot, New Mexico. The ARC-managed facilities at APO consist of a 3.5-m general-purpose telescope and its instruments; and the telescopes, instruments, and ancillary support systems developed for the Sloan Digital Sky Survey (SDSS).

The resources for operating ARC facilities have been provided by participating institutions and private, federal and international sources. The ARC Secretary-Treasurer, Ron Irving, and the ARC Business Manager, Michael Evans, administer the funds received by ARC. ARC disburses some of these funds to the participating institutions through formal agreements and the remainder through contracts managed directly by ARC. New Mexico State University, an ARC member, manages the operations at APO for ARC. The ARC Business Manager directly administers large contracts with vendors when it is advantageous to ARC.

Oversight of ARC operations is the responsibility of the Board of Governors, hereafter the Board. The institutions that constitute the Board are the University of Chicago, the University of Colorado, the Institute for Advanced Study, Johns Hopkins University, New Mexico State University, Princeton University, University of Virginia, and the University of Washington.

Two members represent each institution on the Board and the members are drawn from active scientists and senior university administrators. The Chair of the Board is Professor Edwin L. Turner, from Princeton University. The Board directly oversees the 3.5-m telescope program but created the Advisory Council to advise the Board on matters related to the SDSS-III, in view of the large scale of the SDSS-III program and the large number of institutions participating in that program.

1.2. The Advisory Council

The Board has delegated the oversight and management of SDSS-III to the Advisory Council (AC). The Interim Chair of the Advisory Council is Suzanne Hawley of the University of Washington; a new chair will be elected in the near future.

Institutional membership on the Advisory Council is granted to those ARC institutions that make significant cash or in-kind contributions toward the operation of SDSS-III. Non-ARC institutions become Full or Associate Participants through Memoranda of Understanding (MOU) with ARC. Each Full Participant institution may appoint one voting member to the Advisory Council, and Associate Participant institutions that have joined the survey with at least three named participating scientists may appoint one voting member to the Advisory Council. Participation Groups with 3 or more slots also may appoint a voting member to the AC; those with 10 or more slots may appoint two voting members.

Advisory Council actions are governed by the Principles of Operation-III (PoO-III). The PoO-III has been approved by the Board of Governors and has been agreed to by all institutions joining SDSS-III. Amendments to the PoO-III must be approved by the Board of Governors on the recommendation of the AC by processes described in the PoO-III.

1.2.1. The Steering Committee

Prior to the establishment of the Advisory Committee, the ARC Board of Governors established a Steering Committee to oversee the project and in particular to negotiate with potential institutional partners. The Steering Committee was chaired by Rich Kron. It has now been dissolved, with its role passing to the AC.

1.3. The Central Project Office

1.3.1. The Director

The Board delegates to the Director the executive authority for the operation of SDSS-III. To this end, the Director is responsible for organizing and directing all aspects of the Survey, including the appointment of key personnel. The Director is responsible for ensuring that the available resources are effectively applied toward the scientific goals described in the PoO-III.

The Board, taking into consideration the recommendation of the Advisory Council, appoints the Director for a fixed term. The Board appointed Daniel Eisenstein of The University of Arizona as the Director for an initial 19-month term beginning April 2007. At the November 2008 meeting, the Board appointed Eisenstein to a second term ending in December 2011.

The Director is responsible for preparing the SDSS-III Project Execution. He leads the preparation of funding proposals for the operation of the Survey. ARC submits these proposals to federal agencies and philanthropic institutions and the Director serves as the Principal Investigator for these proposals. He is responsible for drafting, for concurrence by the Advisory Council and approval by the Board, the Memoranda of Understanding with Participating Institutions.

The Director submits both an annual budget and a total budget for the completion of the Survey to the Advisory Council. These budgets include all funds and in-kind services needed for the operation of the Survey, including the acquisition, processing, archiving, and distribution of data to the collaboration and general public. The Advisory Council transmits the Director's budgets along with its recommendations to the Board for approval.

The Director is responsible for implementing financial controls within the project, assisted by the Program Manager and ARC Business Manager. The Director approves all expenditures above $3000. The Director and ARC Business Manager approve all computer purchases in accordance with ARC corporate policy. The Program Manager tracks expenditures against the approved budget and advises the Director of financial status and performance.

The ARC corporate office, under the general supervision of the ARC Secretary/Treasurer, assists the Director and Program Manager with the preparation of the annual budgets and quarterly progress reports. The ARC Business Manager provides quarterly Revenue and Expenditure Reports to the Director, Program Manager, and Advisory Council to show expenditures and obligations compared to the annual budget.

1.3.2. Project Scientist

The Director has appointed David Weinberg of Ohio State University as the Project Scientist. The Director has delegated to the Project Scientist the responsibility for providing the overall quality assurance for the Survey and ensuring its scientific integrity. The Project Scientist monitors the performance of all systems and evaluates the scientific impact of proposed changes to hardware, software, and operating plans. The Project Scientist tracks the progress made on the development of critical new hardware and software systems before they can be certified as meeting the science requirements. He is responsible for assuring the Director that the performance of all systems will meet the scientific goals of the Survey.

The Project Scientist is responsible for definition and documentation of the science requirements for all aspects of SDSS-III. The Project Scientist is responsible for developing and coordinating tests of the scientific integrity of SDSS-III data, comparing the results of these tests with the requirements, and suggesting action to address concerns where these requirements are not being met and the science goals of the project might be compromised.

1.3.3. Program Manager

The Director has appointed Bruce Gillespie of Johns Hopkins University as Program Manager. The Program Manager assists the Director in the performance of his responsibilities. The Program Manager is responsible for developing and maintaining project schedules. He is responsible for preparing annual and cost-to-complete budgets for consideration by the Director and the Project Scientist prior to their submission by the Director to the Advisory Council. He is responsible for tracking project expenditures against the approved budget and reporting them, together with any deviations from the approved budgets, to the Director on a timely basis. He is responsible for preparing annual and quarterly reports that are distributed to the Advisory Council and the funding agencies.

The Program Manager oversees day-to-day operations associated with Survey Operations. He coordinates the engineering effort at APO with the efforts of the engineering groups at the Participating Institutions and the requirements of the observing program. He coordinates the software effort at the Participating Institutions with the requirements of the data processing and distribution programs. He identifies resources at the Participating Institutions when additional resources are needed to meet schedules.

The Program Manager is responsible for informing the Director and the Project Scientist of the state of compliance of Survey Operations with survey metrics. The Program Manager is responsible for developing and maintaining the project schedule and determining the schedule performance of Survey Operations. The Program Manager is responsible for negotiating and managing the institutional work agreements. The Program Manager maintains the list of new work requests until each new project is approved, assigned, and integrated into the overall project plan; or rejected based on the lack of sufficient justification.

1.3.4. Management Committee

The Management Committee provides a forum for the discussion and framing of issues that require action by the Director and/or Advisory Council. The Management Committee is composed of the Director, the Project Scientist, the Program Manager, the Spokesperson, the Data Coordinator, the Survey Coordinator, the Infrastructure Lead, the Observing Systems Lead, the Operations Manager, the Intersurvey Science Coordinator, and the four survey PIs. The Director is the Chair. The Management Committee examines and acts on all issues that have a broad impact on the Survey. When the resolution of an issue requires the approval of the Advisory Council or the Board, the Management Committee reviews the matter and formulates a recommendation for action by the Advisory Council.

The Executive Management Committee (EMC) consists of the Director, Project Scientist, Program Manager, and Spokesperson. The Principles of Operation calls for the EMC to approve External Collaborators.

1.3.5. Hardware and Software Review Boards

Review Boards are responsible for critical assessment of the readiness of hardware and software at each important step in their development.

An Instrument Board chaired by the Program Manager oversees the management of the BOSS spectrograph upgrade, construction of the APOGEE and MARVELS instruments, and related common technical infrastructure improvements. The Instrument Board will conduct technical reviews and monitor the schedule and budget performance of each major hardware task. Release of funds for construction activities will be contingent on successful Board reviews at each stage. Current members of the Instrument Board are Bruce Gillespie (JHU/APO, chair), Jim Gunn (Princeton), Stephen Smee (JHU), French Leger (APO), and Darren Depoy (TAMU).

Software development is reviewed in a similar manner by the Software Board, which is also chaired by the Program Manager. The Software Board will be constituted in the near future.

1.3.6. Change Control Board

As described in section 15, certain aspects of the survey are managed by a formal change control process. The change control board must approve changes to these documents and procedures. The change control board (CCB) consists of the Director, Project Scientist, Observing Systems Lead, Survey Coordinator, Intersurvey Science Coordinator, and Program Manager.

1.3.7. External Advisory Committees

The Director may form ad hoc external Advisory Committees to provide him with advice on the capabilities of the Survey to acquire and process data and to distribute the archived data products to the Collaboration and general astronomy community. These committees will review the effectiveness of observing operations, data processing, and the distribution of data and make recommendations for improvement if survey operations do not meet the expectations of the sponsors. The members of these committees will consist of scientists (primarily astronomers), engineers, and computer professionals with experience in large projects, and who are not engaged in the Survey.

1.4. Survey Teams

The four surveys are the Baryon Oscillation Spectroscopic Survey (BOSS), the Sloan Extension for Galactic Understanding and Exploration 2 (SEGUE-2), the Multi-object APO Radial Velocity Exoplanet Large-area Survey (MARVELS), and the APO Galactic Evolution Experiment (APOGEE). The surveys themselves are described in section 2. Each survey has a survey team that will design and implement the hardware, software, and experimental design necessary to meet the science requirements of the survey. The teams are responsible for the data reduction and quality assurance. They will work with the survey coordinator and operations staff to ensure that plate design and other operations tasks are completed on schedule and meet the survey requirements.

Each team is led by a Principal Investigator (PI). The PI is responsible for the work assigned to the survey team. He or she organizes the survey team, leads the development of budgets and schedules, and monitors and reports performance to the central project office. The PIs are David Schlegel (LBL, BOSS), Connie Rockosi (UC Santa Cruz, SEGUE-2), Jian Ge (Florida, MARVELS), and Steven Majewski (Virginia, APOGEE). The PIs are members of the Management Committee.

Each team has a Survey Scientist (SS) who is responsible for writing the science requirements and target selection documents for the survey and for ensuring that the survey data products meet these requirements. The Survey Scientists report to the Project Scientist. The Survey Scientists are Martin White (LBL/Berkeley, BOSS), Tim Beers (Michigan State Univ., SEGUE-2), Holland Ford (Johns Hopkins University, MARVELS), and Ricardo Schiavon (Gemini, APOGEE).

Each team has an Instrument Scientist (IS) who is responsible for the design, construction, and maintenance of the survey instrument. The Instrument Scientists are Natalie Roe (LBL, BOSS), Connie Rockosi (UC Santa Cruz, SEGUE-2), Xiaoke Wan (Florida, MARVELS), and Michael Skrutskie (Virginia, APOGEE).

BOSS, MARVELS, and APOGEE will have project managers during their hardware development phases. These project managers report to the Program Manager.

1.4.1. Intersurvey Science Coordinator

The Director has created the position of Intersurvey Science Coordinator to advocate for the scientific coordination of the four surveys, including the identification of new science opportunities and the generation of ancillary science target classes. Michael Strauss (Princeton Univ.) will serve in this position and represent these interests on the Management Committee.

1.5. Infrastructure Development

The task of developing the upgrades to the observing system, beyond the immediate scope of the instruments, will be led by the Infrastructure Lead, Jim Gunn (Princeton Univ.). These include upgrades to the cartridge, plug plate, and fiber systems, the operations software, the telescope guider, and the expansion of the support building.

1.6. Survey Operations

The Survey Operations group is responsible for coordinating the observing plans of the four surveys, implementing the resulting unified plan that achieves the strategic and tactical goals of the project, and reporting the observing progress metrics to the Management Committee. The Survey Coordinator, Don Schneider of Pennsylvania State University, leads the survey operations group. The Survey Coordinator will assist all survey teams in designing observing plans that assure that nights are fully utilized, including the constraints of mapping large regions of the sky over multiple years. The Survey Coordinator will oversee the translation of these plans into plug plate designs (which must incorporate the constraints of observing time and sky position to produce hour angle forecasts).

The Survey Coordinator produces the monthly observing plan, which specifies which surveys are intended to be observed each night and lists the observing priorities. Final decisions on what to observe remains with the night observers. The Survey Coordinator will work closely with the Lead Observer to ensure that the observing plans and priorities are clearly transmitted and to ensure that the observing operations group can communicate problems that might affect survey strategy back to the plate design and observation planning process.

The Survey Coordinator will report monthly to the collaboration about the observational progress of the surveys. He will work with the Project Scientist and the Survey Scientists to define quantitative metrics by which the survey progress can be judged. These metrics will be reported quarterly.

The Survey Coordinator will direct the Plate Designer to ensure that plates are designed and fabricated in time for observations. The target lists for the plates are supplied by the Survey Teams. The plate drilling occurs at the University of Washington machine shop under the management of Michael Evans (UW).

1.7. Observatory Operations

Apache Point Observatory is operated by the Astrophysical Research Consortium. ARC has appointed Kurt Anderson (NMSU) as the Site Director and Mark Klaene (APO) as the Site Operations Manager. These two are responsible for all site operations, including those of SDSS-III.

Observatory Operations are carried out under the direction of the Site Director and Site Operations Manager. APO provides all of the basic services and facilities to the technical groups that are needed to carry out their work at the site. APO provides and trains the observing staff that carries out the observations for the Survey. The Site Operations Manager is responsible for providing the observer team with office and laboratory space, onsite and offsite computer networks, and desktop computing. The Site Operations Manager is responsible to ARC for the safe conduct of all activities at the Observatory. The APO Safety Officer, Mark Klaene, provides safety oversight for all activities at APO, establishes the qualifications for all people to engage in various tasks while working at the Observatory, and maintains their training records. In order to fulfill this responsibility, APO provides the safety training for staff engaged in activities at the Observatory.

The Project Scientist, in consultation with the Observing System Lead, is responsible for setting the technical direction and goals of the technical groups and for reviewing the level of observatory support to assure that it is sufficient for the proper execution of the Survey. He sets priorities when the goals of the groups are in conflict.

The Program Manager is responsible for overseeing day-to-day operations of the various project components, planning the general application of resources on an annual basis, and reacting to immediate needs of the survey. The survey operations work plan is coordinated through a weekly teleconference chaired by the Program Manager. The participants in the weekly conference include the Survey Coordinator, Site Operations Manager, Lead Observer or representative, Lead Engineering, and key personnel at the APO site and participating institutions.

The Site Operations Manager will have authority over access of the survey teams to the site. The Lead Engineer will oversee access to the telescope and instruments so as to ensure that commissioning and maintenance activities do not interfere with each other or with the observing program.

Stephanie Snedden (APO) is the Lead Observer; Kaike Pan (APO) is the Deputy Lead Observer. The Lead Observer is responsible for scheduling the observing staff, implementing the monthly observing plans in concert with the Survey Coordinator, ensuring that the observing procedures are well-documented and followed, and training new observers.

1.7.1. Observing Systems

The Observing Systems group is responsible for maintaining the Observing Systems in an operational state throughout the observing phase of operations. The Observing Systems consist of the 2.5-m telescope, the CCD imaging camera, the dual SEGUE-2 and BOSS spectrographs, the MARVELS instrument, the APOGEE instrument, the 2.5-m instrument change system, the equipment for plugging spectroscopic plates, and the data acquisition system for both telescopes. The group is responsible for maintaining and improving the software used to operate the telescopes and instruments (the Observers' Programs) and the data acquisition equipment at APO. In addition, it is responsible for implementing incremental improvements that will increase the efficiency of these subsystems and QA of the raw data. Finally, it is responsible for assuring that the aforementioned systems can meet the Science Requirements.

Observing Systems is led by Jim Gunn (Princeton University) with Connie Rockosi (UC Santa Cruz) as deputy. French Leger is the Lead Engineer; Dan Long is his Deputy. The observing systems group will rely upon the Instrument Scientists with regard to the APOGEE, BOSS, MARVELS, and SEGUE instruments.

1.7.2. Data Archiving

The operations software is responsible for capturing the data stream from the instruments and storing it on disk in a suitable file structure, typically in nightly sets.

The Data Distribution group, under the lead of the Data Coordinator, Michael Blanton of New York University, is responsible for the archiving the data from this point. The Data Archivist, Ben Weaver, will be principally responsible for designing the archiving system.

Raw data will be transferred by internet to the Science Archive Server and the Science Archive Mirror. A copy will also be kept at APO in a safe location. Check-sums will be used to ensure that all network transfers are completed without error.

1.7.3. Role of Survey Teams in Operations

The survey teams undertake efforts that assist survey operations. These efforts include the design of target-selection algorithms, the specification of required calibrations, monitoring pipeline outputs for quality assurance of the data, optimizing the sequence of observations for the end-game, and writing technical papers. The survey teams also serve as centers of expertise to advise on matters of the optimal observing strategy, necessary systems or software development, analysis software, the specific content of periodic data releases, and other matters related to operations.

1.8. Data Processing and Distribution

The Data Distribution group is responsible for all data flow within the project, starting with the nightly data collections produced by the operations software and ending with the public databases. The Data Coordinator, Michael Blanton (NYU), is the lead of this group. The Data Manager, Tom Throwe (BNL) handles the overall scheduling and management of the team. The survey teams are responsible for writing the data reduction pipelines and performing the data reductions, but the schedule for performance of the data reductions will be under the direction of the Data Coordinator.

For reference in this section, the data distribution to the collaboration and the public is performed through two interfaces, the Science Archive Server (SAS) and the Catalog Archive Server (CAS). The SAS provides simple flat file access including images and spectra, while the CAS provides access to catalogs and JPEG images through a SQL database. The SAS serves as the source of raw data for the individual reduction teams. The detailed data distribution plan is described in more detail in Chapter 13.

BOSS imaging, targeting and spectroscopic data reductions are the responsibility of the BOSS PI (David Schlegel), and will be executed by the reduction team at LBL with assistance by the Data Archive Scientist (Benjamin Weaver). The BOSS PI is responsible for procuring adequate computational resources. The computer cluster for BOSS reductions is maintained and run by the LBL Scientific Cluster Support team.

SEGUE-2 targeting and spectroscopic data reductions are the responsibility of the SEGUE-2 PI (Connie Rockosi), and will be executed by the reduction team at Princeton (Steve Bickerton and Fergal Mullally). The SEGUE-2 PI is responsible for procuring adequate computational resources. The computer cluster for SEGUE-2 reductions is maintained by the Princeton systems administrators.

MARVELS targeting and spectroscopic data reductions are the responsibility of the MARVELS PI (Jian Ge), and will be executed by the reduction team at the University of Florida (currently Suvrath Muhadevan, Sivarani Thirupathi, Brian Lee, Scott Fleming). The MARVELS PI is responsible for procuring adequate computational resources. The computer cluster for MARVELS reductions is maintained by a University of Florida systems administrator.

APOGEE targeting and spectroscopic data reductions are the responsibility of the APOGEE PI (Steven Majewski), and will be executed by the reduction team at the University of Virginia. The APOGEE PI is responsible for procuring adequate computational resources. The computer cluster for APOGEE reductions is maintained by the UVa systems administrators.

Each survey's PI will appoint a reduction team member to act as a liaison to assist with validation of the data assembly step and with archiving and documentation of their surveys' reduced data sets.

The Data Coordinator has overall responsibility for the archiving and distribution of the raw and reduced data. The Data Archive Scientist at NYU will organize and maintain of the Science Archive Server (SAS) and Science Archive Mirror (SAM). He will liaise with each reduction team to organize transfer of raw and reduced data to and from the SAS. The Data Manager will handle procurement of adequate computational resources at SAS and SAM. The SAS will exist at LBL, on clusters maintained by the LBL Scientific Cluster Support team. The SAM will exist at NYU, on clusters maintained by the Center for Cosmology and Particle Physics.

Once the reduced data from the reduction teams are uploaded to the SAS, the data must be packaged into an integrated scientific form, a step we will refer to as “data assembly.” The Data Coordinator will execute these data assembly steps, with the assistance of the Data Archive Scientist, the Drilling Coordinator, and the survey team liaisons where appropriate.

The Database Development Group (DDG; currently led by Alex Szalay and Ani Thakar) is responsible for the organization and maintenance of the Catalog Archive Server (CAS). The CAS will exist at JHU, on clusters maintained by the CAS team systems administrators. A database administrator will maintain the CAS. The Data Manager will handle procurement of adequate computational resources for CAS. CAS mirrors will be maintained at several offsite locations, with local administrators maintaining them.

The Data Coordinator is responsible for the database loading process. Under the direction of the Data Coordinator and Data Manager, the reduction team liaisons will be responsible for producing SQL-format schema and documentation. The Data Coordinator will execute the transformation of standard format files to database-loading (CSV) files. The DDG head and the Data Archive Scientist will execute the actual transfer process.

Weekly teleconferences chaired by the Data Coordinator review the state of data processing and distribution, discuss the results of ongoing tests of processed data, and set priorities for upcoming effort.

SDSS-III Participating Institutions may wish to establish informal web sites to host SDSS-III data, but ARC takes no responsibility for these sites.

1.9. Collaboration Affairs

The Collaboration consists of scientists from the Participating Institutions. Its membership also includes scientists from non-participating institutions who have earned data access rights through their contributions to survey infrastructure, and external collaborators who provide expertise on specific projects. The Collaboration provides opportunities for its members to exchange information and ideas freely, thereby assisting the pursuit of their individual research goals. Collaboration Affairs was formed to accomplish these goals.

1.9.1. Scientific Spokesperson

The Director has delegated the management of the affairs of the SDSS-III collaboration to the Scientific Spokesperson. He has also delegated to the Spokesperson the primary responsibilities for representing the Survey to the scientific community and for raising the visibility of the Survey within the astronomy and physics communities. Accordingly, the Spokesperson oversees the publication of scientific, technical, and data release papers. The Spokesperson is also charged with promoting the collective scientific productivity of the SDSS-III Collaboration.

The Spokesperson is elected by a majority vote of the Participants for a three-year term. In the event the Spokesperson resigns or becomes unable to serve, the Director will appoint an acting Spokesperson until a new election can be held.

The Spokesperson is responsible for organizing Collaboration Affairs, chairing the Collaboration Council, and nominating the Technical Publications Coordinator and Scientific Publications Coordinator. A major responsibility of the Spokesperson is to create a healthy collegial environment in which the pursuit of the scientific goals of the surveys can flourish. Special attention is paid to the mentoring of postdocs and graduate students, and to rules governing graduate student theses involving Survey data. The principles guiding the work of the Collaboration are spelled out in the PoO, Publications Policy, and similar documents. As required, the Spokesperson is responsible for proposing and/or revising collaboration-specific policies for approval by the Management Committee and the Advisory Council.

The Spokesperson arranges for the organization of presentations at the meetings of professional societies, in the course of discharging his responsibilities for representing the Survey to the scientific community and for raising the visibility of the Survey within the astronomy and physics communities. He consults CoCo on these matters and brings them to the attention of the Management Committee. The Publications Office provides the Collaboration with a means to oversee the preparation of technical and scientific publications.

The Spokesperson, with the help of CoCo, solicits offers from Participating Institutions to hold annual Collaboration meetings. The organization for each meeting is the responsibility of the local organizing committee. The agenda and special events are designed and approved by the Spokesperson and CoCo.

1.9.2. Collaboration Council

The Collaboration Council (CoCo) assists the Spokesperson in the management of Collaboration Affairs. It provides advice to the Spokesperson on all Collaboration matters, including recommendations for policies on publications, scientific representation, and science projects. The membership consists of one person from each Participating Institution or Participation Group that has a vote on the Advisory Council. The CoCo representative is appointed by the member(s) of the Advisory Council from that institution. In addition, one representative is elected by the non-voting Associate Members and External Participants to serve on CoCo. The Spokesperson serves as the Chair of CoCo.

The Chair of CoCo is responsible for ensuring that CoCo meets regularly via teleconference to discuss matters pertaining to the health of the Collaboration and to advise the Spokesperson.

CoCo is responsible for reviewing proposals for science projects that include external collaborators. External collaborators bring special expertise to a particular project and the capacity to enable the project so that its results can be published in a timely way. In exchange for their assistance, external collaborators are given access to the data appropriate for the specific project and become eligible for authorship on that research. If the CoCo determines that a proposed external collaboration is to the overall benefit of the SDSS-III Collaboration, they will recommend approval to the Executive Management Committee, who reviews such collaborations in light of project fund-raising and other goals.

CoCo is responsible for organizing and conducting the election of the Spokesperson.

1.9.3. Education and Public Outreach

The principal goal of the SDSS-III Education and Public Outreach (EPO) effort is to make the discoveries, data, and methods of the SDSS-III intelligible and interesting to a broad audience of non-scientists. The SDSS-III EPO Coordinator will be Jordan Raddick of Johns Hopkins University. He will review the current suite of EPO efforts already underway and assemble them in creative ways, identifying strengths around the Collaboration that may be put together for an even greater benefit.

The Public Information Officer handles press releases and communication with the news media. The Spokesperson is responsible for providing the Public Information Officer with the scientific and technical information that will be distributed in press releases and other communications with the media. The Public Information Officer is responsible for coordinating and organizing the work of the Public Information Officers at the other Participating Institutions in order to assure that the interests of all of the Survey sponsors are properly served. Jordan Raddick of Johns Hopkins University has agreed to serve in this capacity.

Jordan Raddick will also manage the SDSS-III public web site, as this activity overlaps well with the EPO and press activities.

1.9.4. Publications Office

The Spokesperson oversees the Publications Office, which provides a means to disseminate scientific results in draft form to the collaboration. This enables review of the scientific content of draft papers.

The Scientific Publications Coordinator and the Technical Publications Coordinator are appointed for the duration of the SDSS-III survey by the Spokesperson. The Spokesperson maintains a web page on www.sdss3.org listing all published papers and papers approved for publication in refereed journals. These papers include papers posted on astro-ph and conference proceedings. A separate, internal web page, accessible only to the Collaboration, is maintained for work in progress prior to its acceptance for publication. Policies and procedures for the publications are posted on www.sdss3.org.

1.9.5. Ombudsman

The ARC Board will appoint a standing Ombudsman for the project to help resolve disagreements arising in any aspect of the project in an informal manner. As a neutral third party, the Ombudsman does not advocate for the project or for either party in a dispute. The objective is to provide a process for achieving a fair and reasonable settlement working within existing policies and procedures. When a request for services is received, the Ombudsman will work with each party to identify appropriate alternatives that address the conflict and to achieve a mutually satisfactory resolution. Consultation with the Ombudsman does not preclude later pursuit of a resolution through formal channels if that is still desired.

2. Baseline Project Definition

2.1. Overview

SDSS-III is a six-year program (summer 2008 – summer 2014) of four surveys on three scientific themes: dark energy and cosmological parameters, the structure and history of the Milky Way, and the population of giant planets around other stars. The four surveys are summarized briefly here and described more completely in the section 2.3. All surveys are conducted using the Sloan Foundation 2.5-meter telescope at Apache Point Observatory.

- **The Baryon Oscillation Spectroscopic Survey (BOSS):** BOSS will measure spectroscopic redshifts of 1.3 million luminous galaxies to a limiting redshift $z \approx 0.75$ and spectra of 160,000 QSOs at z=2-3.5, to enable measurement of the baryon acoustic oscillation (BAO) scale in the galaxy distribution and the Lyα forest. These measurements will provide strong new constraints on the equation of state of dark energy and the curvature of space. BOSS will also provide one of the most powerful data sets for studying other aspects of large-scale structure at these redshifts and for investigating the evolution of massive galaxies and high-redshift quasars.
- **The Sloan Extension for Galactic Understanding and Exploration 2 (SEGUE**-2): SEGUE-2 will obtain optical spectra of 140,000 stars in the outer Galaxy, with resolution R=1800 and typical signal-to-noise ratio S/N = 20-25 per pixel, yielding typical velocity precision of 5-10 km s^{-1} and abundance precision of 0.2 dex in [Fe/H]. These data, in combination with 240,000 similar spectra from SEGUE, will be used to map the kinematic structure and chemical abundance distribution of the Galactic halo and thick disk and to identify rare classes of stars.
- **The Apache Point Observatory Galactic Evolution Experiment (APOGEE):** APOGEE will obtain *H*-band spectra (1.51-1.65μm) of 100,000 red giant stars from all regions of the Galaxy, with resolution R=22,500-25,000 and typical signal-to-noise ratio S/N=100 per pixel, yielding typical velocity precision of 0.5 km s^{-1} or better and allowing individual abundance determinations of ~15 chemical elements in each program star. These data will revolutionize the study of the kinematical structure and chemical evolution of the Galaxy.
- **The Multi-object APO Radial Velocity Exoplanet Survey (MARVELS):** MARVELS will monitor the radial velocities of ~11,000 program stars with visual magnitude V=8-12, visiting each star ~30 times over an 18-month interval. The velocity precision will be approximately 14, 18, and 35 m s^{-1} at V=8, 10, 12, including systematic errors. MARVELS will provide the comprehensive and well characterized statistical data set needed to test predictive theories of the formation and dynamical evolution of giant planets.

NSF funds support only BOSS and APOGEE, and DOE funds support only BOSS. The prescription for assigning costs to surveys is described in Section 20.6.

SEGUE-2 uses the existing SDSS spectrographs. BOSS will use a significantly upgraded version of these spectrographs with more sensitive detectors, higher throughput optics, and 1000 fibers per spectroscopic plate (compared to the current 640). APOGEE will use a new 300-fiber, cryogenic *H*-band spectrograph. MARVELS uses a 60-fiber interferometric spectrograph called ET1 (Exoplanet Tracker 1), which has been provided to SDSS-III as an in-kind contribution from the University of Florida. The full program described above assumes that a second 60-fiber spectrograph, ET2, will be added for the final four years of MARVELS. However, the construction of ET2 is beyond the scope of the baseline SDSS-III budget and will require separate funds. Construction of the BOSS spectrograph upgrades and the APOGEE spectrograph is part of the baseline SDSS-III budget.

SEGUE-2 will be the principal dark-time survey until the BOSS spectrographs are commissioned in fall 2009. BOSS will also carry out dark-time imaging in the southern Galactic cap during fall 2008. After the BOSS spectrographs are commissioned, BOSS will be the sole dark-time survey for the remaining five years of SDSS-III. MARVELS will be the sole bright-time survey until the APOGEE spectrographs are commissioned in winter 2011. Thereafter, APOGEE and MARVELS will share the bright time equally. For the most part, APOGEE and MARVELS will carry out simultaneous observations of the same fields, sharing the focal plane, so that each survey will be able to use close to 100% of the available bright time. An additional bright-time program, SEGUE-bright, would use the BOSS spectrographs to obtain optical spectra of bright stars, sharing the focal plane with APOGEE and MARVELS. The cost of SEGUE-bright, while relatively small (~$900K), is beyond the scope of the baseline SDSS-III budget.

2.2. Science Requirements and Target Selection Documents

Each of the four surveys is described by a Science Requirements Document, which defines the high-level goals of the survey and the implied requirements on instrument and software performance. This document is developed by the Survey Scientist and the Survey PI, in consultation with other members of the survey team and with the SDSS-III Project Scientist, and brought forward to the Central Project Office for review and approval.

Once the requirements document is accepted, changes must be approved by the Change Control Board. If the anticipated instrument or software performance will not meet the science requirements, the Director, in consultation with the survey team, the Central Project Office, and, if necessary, the Collaboration Council and/or the Advisory Council, will recommend either a revision of the requirements or a plan (which may include additional budget allocation) for improving performance to meet the requirements.

Each survey also provides a Target Selection Document, which defines the methods and criteria used to select spectroscopic targets. This document follows the same development and acceptance procedures as the Survey Requirements Document. Once the Target Selection Document is accepted, changes to the survey target selection must be approved by the Change Control Board.

2.3. Survey Goals: Science Outcomes and Data Products

Each of the four surveys has the goal of producing data sets with well defined properties, designed to achieve a set of science outcomes. This section, derived from the survey Science Requirements Documents, summarizes these data products and science goals. The data products will be made public through the Science Archive (but see the caveat below about MARVELS). The science outcomes will be achieved through analysis of these data products by members of the SDSS-III Collaboration and by others in the astronomical community.

2.3.1. BOSS

The defining goals of BOSS are to use BAO as a standard ruler to measure the cosmic distance scale and cosmic expansion history and thereby probe the causes of cosmic acceleration. More specifically, the goals of the LRG (Luminous Red Galaxy) survey are to measure the absolute angular diameter distance, d_A, to z = 0.35 and z = 0.6 to 1.0% and 1.1% precision, respectively, and to measure the absolute Hubble parameter, $H(z)$, at z = 0.35 and z = 0.6 to 1.8% and 1.7%, respectively. The goals of the QSO Lyα forest survey (hereafter referred to as the QSO survey) are to measure d_A and $H(z)$ to precision of at least 1.7% and 1.5%, respectively, at an average redshift z=2.5.

The BAO distance measurements and other measurements of large-scale structure in the LRG and QSO surveys can be used to constrain cosmological parameters, and a precision measurement of a number of these parameters is an additional goal of BOSS. The forecast precision of the constraints depends on additional assumptions about the model space and the constraints from complementary experiments. For the model space and forecasts of "Stage II" (current generation) dark energy experiments adopted by the Dark Energy Task Force, the forecast constraints from BOSS BAO measurements include determinations of the equation-of-state parameter w (at pivot redshift $z \approx 0.25$), space curvature Ω_K, and Hubble constant h with 1σ uncertainty $\Delta w = 0.029$, $\Delta\Omega_K = 0.002$, and $\Delta h = 0.008$. The precision of these determinations improves considerably if one includes the forecast constraints from BOSS broad-band power and redshift-space distortion measurements and/or from other "Stage III" dark energy experiments.

In addition to large-scale structure and cosmological parameter measurements, BOSS aims to provide a powerful database for studying the evolution of massive galaxies at $z < 0.7$ and the luminosity function and clustering of quasars at $2.3 < z < 6.5$. These goals do not drive technical requirements, but the survey we have designed to achieve precision distance measurements will also provide a spectroscopic data set of unprecedented size for both populations of objects. Design choices (e.g., on target selection) that are roughly neutral with regard to cost and impact on the distance measurement goals are guided by these additional science considerations.

The principal BOSS data products will be:

- Calibrated spectra of 1.3 million luminous galaxies selected from SDSS *ugriz* imaging over 10,000 square degrees. The galaxy sample will have an approximately constant space density of $n=3\times10^{-4}\ h^3\ \mathrm{Mpc}^{-3}$ to z=0.6, thereafter falling linearly out to a limiting redshift z=0.75-0.8, with limiting apparent magnitude $i\approx19.8$. The spectral resolution of BOSS spectra will be in the range R=1500-2500, with some dependence on wavelength. The expected rms redshift precision for the faintest galaxies is 70 km s^{-1} (though the requirement is only 300 km s^{-1}), and for most galaxies the spectra will be good enough to provide additional information about the age and metallicity of the stellar population.
- Calibrated spectra of 160,000 QSOs with z>2.2, selected from SDSS *ugriz* imaging with the resolution (see above) and S/N required to measure the flux power spectrum of the Lyα forest accurately on BAO scales. These spectra will also provide information about the redshifts and properties of the QSOs themselves, and they can be used to create large samples of metal-line absorption systems and high column density hydrogen absorbers.
- Catalogs of galaxies and QSOs with redshifts and other properties derived from these spectra and from SDSS imaging.
- New imaging of approximately 2000 square degrees in the southern Galactic cap, to increase the total volume of BOSS and enable us to make efficient use of fall spectroscopic observing time. The science requirement is a contiguous footprint of 2200 deg^2, including the approximately 750 deg^2 of pre-existing imaging from SDSS-I.
- BOSS will also devote up to five percent of its fibers to ancillary science targets, with surface densities up to 10 deg^{-2}, allowing large samples of interesting but rare objects.

2.3.2. SEGUE-2

The defining goal of SEGUE-2 is to map stellar populations and their kinematics over a large volume of the Galaxy, from the inner halo to large Galactocentric distances where late-time accretion events are expected to dominate the halo population. SEGUE-2 data, together with data (now public) from the SEGUE component of SDSS-II, will constrain existing models for the formation of the stellar halo and inform the next generation of high resolution simulations of Galaxy formation. A large fraction of SEGUE-2 fibers are devoted to faint targets that are likely to be high luminosity and therefore distant tracers of the halo: metal-poor main sequence turnoff stars, blue horizontal branch stars, and K-giants. A sample of candidate M-giant stars will provide additional probes of the outer halo. Other SEGUE-2 target classes will contribute to our understanding of old Galactic stellar populations and provide important samples for testing aspects of stellar astrophysics: candidate very metal poor stars, white dwarfs, and proper-motion selected extreme cool subdwarfs and cool white dwarfs. A sample of candidate high-velocity stars and hyper-velocity stars will probe the physical origin of these extraordinary populations and the shape of the Galactic gravitational potential.

The SDSS-II SEGUE survey obtained "bright" and "faint" plates along each targeted line of sight, along with additional imaging outside the original SDSS footprint. SEGUE-2 does not involve additional imaging, and it will obtain only "faint" spectroscopic plates so that it can better survey the distant halo.

The principal SEGUE-2 data products will be:

- Calibrated spectra of 140,000 stars in the categories listed above, at the same spectral resolution ($R \approx 1800$) as current SDSS and SEGUE spectra. The depth of exposures will be sufficient to yield median S/N=10 per pixel at r=19.5 for metal-poor main sequence turnoff stars. This signal-to-noise ratio is required for successful parameter estimation by the SEGUE stellar parameters pipeline (SSPP). Magnitude limits for other target categories that require parameter estimation are chosen so that they also yield S/N=10 per pixel or greater in the same exposures. In average conditions, the exposure time required to reach this depth is 3 hours. For median targets, the typical S/N will be 25-30 per pixel.
- Spectroscopically derived parameters for these stars, including: radial velocities, effective temperatures, surface gravities, and heavy element abundances ([Fe/H]). Typical errors on these quantities (statistical and systematic based on comparison to high resolution data) are 5-10 km s^{-1}, 140K, 0.29 dex, and 0.24 dex.

In the later years of the survey, we hope to carry out additional, bright-time observations of approximately 100,000 stars to a magnitude limit r=17.5 using the BOSS spectrographs, sharing the focal plane with MARVELS and APOGEE. However, this program, SEGUE-bright, is beyond the scope of the baseline SDSS-III budget. We refer to it briefly elsewhere in the document, primarily in Chapter 9.

2.3.3. APOGEE

The defining goal of APOGEE is to carry out a large scale, systematic, uniform, high-precision spectroscopic survey of *all* Galactic stellar populations (bulge, disk, bar, halo), which can be used: to derive element-by-element abundance data for constraining detailed models of Galactic chemical evolution; to obtain high-precision kinematical measurements for constraining dynamical models of the disk, the bulge, the bar, and the halo, and for identifying substructures in these components; to infer properties of Population III stars, by detecting them directly if they survive to the present day, or by measuring their nucleosynthetic products in the most metal-poor stars that do survive. A critical distinguishing feature of APOGEE, made possible by observing at infrared rather than optical wavelengths and thus suffering vastly less extinction by interstellar dust ($A_H/A_V = 1/6$), is that it will use the same methods and stellar target types to probe all regions of the Galaxy. At the same time, APOGEE will increase the total number of high-resolution, high S/N stellar spectra by two to three orders of magnitude relative to the number available today.

The principal APOGEE data products will be:

- Wavelength-calibrated, sky-subtracted, telluric-absorption corrected, 1-D spectra of 100,000 or more red giant stars primarily selected from the *Two-Micron All Sky Survey* (2MASS), along sightlines that probe all the major components of the galaxy. The spectra will cover most of the wavelength range 1.52-1.68μm, with resolution R=22,500-25,000 and typical signal-to-noise ratio S/N=100 per pixel. Typical observations will have a cumulative exposure time of three hours and reach this S/N at a limiting magnitude *H*=12.5, but some fields will have longer exposures and reach to *H*=13.5 or more.
- Effective temperatures and surface gravities derived from these spectra.
- Radial velocities derived from these spectra, with typical accuracy of 0.5 km s^{-1} or better.
- Abundance measurements for a large number of key chemical elements, including CNO, α-elements, iron-peak elements, and odd-*Z* elements, which provide complementary information about chemical evolution and nucleosynthesis. The "top priority" elements are C, N, O, Mg, Al, Si, Ca, Fe, and Ni, while the "medium priority" elements are Na, S, Ti, Mn, and K; the practical impact of these priority designations is specified in the Science Requirements Document. Lower priority elements are V, Cr, Co; these will be obtained for many stars but do not drive requirements.

2.3.4. MARVELS

The defining goal of MARVELS is to produce a large, statistically well defined sample of ~150 giant planets with periods ≤ 480 days that are drawn from a large sample of host stars that have well understood selection biases and encompass a wide range of stellar properties. MARVELS will reveal the diversity of giant exoplanet populations and provide the premier next-generation data set for testing models of the formation, migration, and dynamical evolution of giant planet systems. Additional goals are to find rare planetary systems (e.g., Very Hot Jupiters, short-period super-massive planets; short-period eccentric planets; highly eccentric planets; planets orbiting low metallicity host stars), to find planets around (sub-) giant stars and active stars, to find transiting planets, to quantify the emptiness of the brown dwarf desert, to find rapidly interacting multiple planet systems, and to identify systems that can be observed at higher RV (radial velocity) precision and/or over longer timescales to find multiple planet systems.

MARVELS will monitor the radial velocities of 10,000 main sequence stars and 1,000 (sub)-giant stars with visual magnitude V=8-12, visiting each star ~30 times over an 18-month interval.

The principal MARVELS data products will be:

- Radial velocity curves of all program stars, with RV precision per observation of 14.3 m s^{-1} (V=8), 15.5 m s^{-1} (V=9), 18.3 m s^{-1} (V=10), 23.9 m s^{-1} (V=11), and 35.0 m s^{-1} (V=12). The quoted precisions include all systematic contributions to the measurement error; the photon-noise limited precision at these magnitudes will be 3.4, 5.4, 8.5, 13.5, and 21.3 m s^{-1}, respectively.
- Characterization (periods, eccentricities, and velocity amplitudes) of the planetary systems identified from these radial velocity curves. Forecasts based on the known statistics of planetary systems indicate that these data will yield approximately 150 new planet discoveries.
- Optical spectra (from the SDSS or BOSS spectrographs) of the program stars.
- Stellar parameters derived from these spectra by the SEGUE stellar parameters pipeline.

There are two important caveats to the above description of MARVELS and its data products. First, the size of the sample is predicated on a second 60-fiber spectrograph, ET2, operating for the final four years of the survey, with the same performance as the instrument (ET1) that has already been provided as an in-kind contribution to SDSS-III from the University of Florida. Funding for this second instrument is not in the scope of the current SDSS-III budget, and if funding for the second instrument cannot be obtained, then the sample of targeted stars and detected planets will be correspondingly smaller.

Second, the scope of the current SDSS-III MARVELS budget is to obtain six years of MARVELS data and to carry out minimal reductions needed to assure that the data are of sufficient quality to meet the science goals. Refined analyses of the MARVELS data and publication of the radial velocity curves and planet detections will be carried out by the MARVELS team on a best-effort basis rather than a required basis, unless additional funds can be raised to restore the scope cuts necessitated by the absence of NSF funding for MARVELS.

2.4. Schedule

The observing systems at Apache Point Observatory will be operated full time from 15 July 2008 through 30 June 2014 for SDSS-III data collection. Some of this time is used for engineering and instrument commissioning, as necessary, and we require down time in the summer for aluminization and other critical maintenance. BOSS and SEGUE-2 observations are carried out only during dark time; MARVELS and APOGEE observations are carried out only during bright time. The definition of dark and bright time is specified in the next section, but roughly speaking bright time is when moonlight makes observations of the faint BOSS/SEGUE-2 targets highly inefficient.

MARVELS pre-selection observations (observations using the SDSS spectrographs to cull unsuitable targets from the preliminary target list) began in July 2008. MARVELS commissioning and survey observations, and additional pre-selection observations, began in September 2008. MARVELS survey observations, and additional pre-selection observations when needed, will continue in bright time through June 2014.

SEGUE-2 observations began in September 2008 and will continue through the 2009 summer shutdown. This will be the end of SEGUE-2 dark time observations unless delivery of the BOSS spectrograph upgrades is delayed, in which case SEGUE-2 will continue in Fall 2009.

BOSS imaging is being carried out during photometric conditions with good seeing in Fall 2008. If the imaging area has not reached the science requirement and the imager remains healthy, additional BOSS imaging may be carried out in Fall 2009.

BOSS spectroscopic observations will begin in Fall 2009, following commissioning of the upgraded spectrographs. They will continue through June 2014.

APOGEE observations will begin in Q2 of 2011, following commissioning of the APOGEE spectrograph. APOGEE and MARVELS will share the bright time equally thereafter. For the most part, the two surveys will carry out simultaneous observations of the same fields, sharing the focal plane, so that each can use close to 100% of the available bright time.

Data are processed and made available to the SDSS-III Collaboration as soon as they are calibrated (Data Archive Server) and as soon as they are loaded into the database (Catalog Archive Server). Loads of the CAS will be carried out at least annually. The schedule of public data releases is given in Chapter 13 below.

In the fourth quarter of each calendar year, an expense summary will be prepared for that year and a budget request for the following year will be presented to the Advisory Council for approval, and subsequent approval by the ARC Board of Governors. Quarterly and annual reports will be prepared for the Advisory Council, the Sloan Foundation, the NSF, and the DOE.

Schedules of the development projects (hardware and software) are discussed in the individual development chapters below, and a detailed comprehensive schedule is presented in Appendix D..

2.5. Lunation and Survey Priorities

The principal division of observing time in any given season is between bright-time and dark-time surveys. "Bright time" is time during which the moonlit sky is too bright for efficient observation of the faint BOSS or SEGUE-2 targets. A single night may be divided between dark and bright time.

More precisely, we define dark time to be time when the moon is below the horizon or at lunar phase > 110 degrees, grey time to be when the moon is above the horizon and at phase between 80 and 110 degrees, and bright time to be when the moon is above the horizon and at phase less than 80 degrees. SEGUE-2 will operate in dark+grey time, and BOSS will operate in dark time only during the fall (when only the south Galactic cap is available) and in dark+grey time during other seasons. MARVELS and APOGEE will operate in bright time, and also in grey time during fall after 2008. With our adopted definition, approximately 60% of observing time is dark or grey and 40% is bright.

During Fall 2008, BOSS imaging has priority when conditions are photometric and of sufficiently good seeing to achieve the necessary depth and photometric accuracy. If BOSS imaging continues in Fall 2009, it will have priority over BOSS spectroscopy when conditions are photometric and seeing is good, until the required contiguous imaging area is achieved. Other than BOSS imaging, all dark time observations are devoted to SEGUE-2 spectroscopy (through summer 2009) and to BOSS spectroscopy (thereafter).

Initially, all bright time will be devoted to MARVELS observations. Once the APOGEE spectrograph is successfully commissioned, the two surveys will operate simultaneously, sharing the focal plane, so that each uses close to 100% of the available bright time. Representatives of the two survey teams will work together with the Director, Project Scientist, Survey Coordinator, and Spokesperson to craft an observing strategy that effectively serves the interests of both surveys. This strategy will incorporate a mix of "long fields" that are observed for ~20-30 hour-long exposures (and thus have roughly the standard MARVELS cadence) and "short fields" that are observed fewer times (typically 3-5 hours of exposure). The long fields will be in areas of high stellar surface density, enabling APOGEE to target multiple independent sets of stars to its standard magnitude limit; in addition, APOGEE will target some stars in these fields to deeper magnitude limits, acquiring a larger total exposure time. The short fields will enable APOGEE to cover sightlines that are important for Galactic structure investigations. In these fields, MARVELS will conduct a "quick-look" survey in which several observations allow identification of radial-velocity variables for follow-up observations with other telescopes.

The Survey Coordinator is responsible for developing monthly observing plans that are consistent with these priorities, and the Lead Observer is responsible for implementing these monthly plans at APO (see chapter 10). The observing plans specify the division of each night between dark-time and bright-time programs. Observers may make decisions during the night that deviate from these plans if they believe they can increase the overall data collection efficiency by doing so. The Survey Coordinator will keep track of time shifted among surveys by these decisions and rebalance if substantial cumulative deviations arise.

If sufficient funds are raised to support the SEGUE-bright observing program, it will share the focal plane with MARVELS and (eventually) APOGEE. Field selection will be driven by MARVELS and APOGEE considerations, with SEGUE-bright considerations coming into play when the other surveys are effectively neutral with regard to field choice. If technical issues make it infeasible to collect data for all bright time surveys simultaneously on a given night (e.g., because of engineering or plate plugging constraints), MARVELS and APOGEE will have priority over SEGUE-bright.

2.6. Achievable Metrics

Each survey has made forecasts of the amount of survey-quality data that can be collected within its allocation of observing time. These baseline forecasts are based on weather statistics obtained during SDSS-I and II.

The metric for BOSS imaging is unique, new sky area observed within the BOSS footprint with data quality sufficient for BOSS spectroscopic target selection. Our baseline forecast indicates that we can obtain the required 1500 deg^2 of new imaging in the Fall 2008 and Fall 2009 seasons. However, the fluctuations of weather make this forecast highly uncertain, as we require photometric conditions with good seeing at the times that the relevant sky area is observable. As of December 15, 2008, the amount of additional BOSS imaging obtained under good conditions is approximately 1300 deg^2.

The metric for BOSS spectroscopy is the number of spectroscopic plates observed (and meeting the requirements for survey quality data). These will be tracked separately for the northern and southern Galactic cap regions. The number of spectroscopic plates in these two regions is approximately 1580 and 520. BOSS requires 2100 plates in total to meet its survey goals. The baseline forecast presented in Appendix A indicates 2% and 45% margins for northern and southern plates, respectively. We hope to increase the BOSS southern imaging area beyond the 2200 deg^2 requirement so that the margin for achieving the total required number of BOSS plates is larger.

The metric for SEGUE-2 spectroscopy is the number of spectroscopic plates observed (and meeting the requirements for survey quality data). SEGUE-2 requires 240 plates to meet its survey goals. Our baseline forecast indicates that we should be able to obtain these SEGUE-2 plates before the summer shutdown in 2009, with a slower rate in Fall 2008 because the best conditions will be devoted to BOSS imaging. We will measure the rate of SEGUE-2 plate acquisition against this baseline.

The metric for MARVELS spectroscopy is the number of target stars observed with radial velocity accuracy meeting the survey requirement, and the number of visits per star. Under the full ET1+ET2 plan, our initial forecasts (based on detailed simulations) indicated that we could obtain an average of 30 observations over an 18-month span for each of 11,000 program stars. The baseline forecast in Appendix A adopts more conservative assumptions about the fraction of clear weather (50% instead of 60%) and about the overhead (22 minutes per 50-minute exposure instead of 11 minutes), and it therefore yields a somewhat lower number of stars. In practice, MARVELS breaks naturally into three 2-year cycles, and we will measure the progress against baseline separately in these three cycles. If ET2 is not implemented in time for the second cycle, or if its capabilities are not identical to those of ET1 (e.g., because it has more or fewer fibers), then we will revise the baseline forecast for the second and third cycles accordingly. We note that the performance metric and forecast do not currently include the "quick-look" MARVELS observations in the "short field" MARVELS/APOGEE co-observations. Once the co-observation strategy is fully defined, we will define a performance metric that includes these observations.

The metric for APOGEE spectroscopy is the effective number of program stars observed at a S/N ratio that meets the survey requirement. By "effective number" we mean the sum of all survey target exposure times divided by three hours, which is the time required to reach H=12.5 at S/N=100 per pixel according to our baseline instrument design --- i.e., it is the number of stars that would be observed if the cumulative exposure time for all stars was three hours. In practice, some brighter stars will be observed for less time and some fainter stars will be observed for more time. Our forecast assuming 50% clear weather, 1.5 hours of overhead per three-hour cumulative exposure, and use of 95% of the allocated APOGEE/MARVELS time (we allow 5% for calibration observations) indicates that we should obtain 104,300 effective targets, a 4.3% margin on the APOGEE goal of 100,000.

Appendix A presents a baseline plan of anticipated progress measured by these metrics for the four surveys. This baseline is based on 50% clear weather (slightly conservative relative to statistics from SDSS-I/II) and our current expectations of exposure times and overhead (again slightly conservative). The Survey Coordinator will track progress of the four surveys against this baseline plan and give a monthly progress report (verbally and by e-mail) to the Management Committee. We will also assess the degree to which progress against the baseline plan is affected (negatively and positively) by weather relative to that assumed in the forecast, so that we can quickly be alerted to hardware, software, or observational problems that cause us to fall behind schedule.

Progress against baselines will be reported in the quarterly reports prepared for the Advisory Council and the funding agencies, and a more complete assessment of survey progress will be included in third-quarter reports (based on progress up to summer shutdown). In the event that a survey is falling consistently behind its baseline, the Director and Project Scientist will formulate a written response to the Management Committee and Advisory Council, recommending either a change in survey operations, a revision of the baseline (including discussion of the scientific impact of this revision), or both. This response will be summarized in the relevant quarterly report.

3. Work Breakdown Structure

The SDSS-III Work Breakdown Structure (WBS) is organized to identify the effort, deliverables, and services necessary to meet project objectives. It is also used to identify organizational and individual responsibilities and integrate project scope, cost, and schedule. Table 3.1 presents a top-level view of the WBS. Appendix B contains a more detailed version.

The WBS organizes work and cost into eight areas: Central Project Office, Infrastructure and Interfaces, SEGUE-2, MARVELS, BOSS, APOGEE, Survey Observing, and Data Distribution. Subsequent sections describe the work and costs captured in each area.

3.1. Central Project Office

1.1 Program Management: Captures the work associated with the overall management of the SDSS-III project, including schedule & budget development, project reviews, reporting, and formation of the collaboration.

1.2 Business Management: Captures the work associated with the ARC business office, including all accounting and account management.

1.3 ARC Affairs: Captures the work associated with ARC membership, such as insurance and annual meetings.

1.4 Survey Documentation: Captures the work associated with documenting the survey practices.

1.5 Collaboration Affairs: Captures the work associated with the administration of the science collaboration, including support of the Spokesperson.

1.5.1 Education and Public Outreach: Captures the work associated with the Education and Public Outreach program.

1.5.2 Public Information: Captures the work associated with press activities and the Public Information Officer.

1.6 Survey Coordination: Captures the work associated with coordinating the operations of the four surveys and the interface to the mountain observing staff.

3.2. Infrastructure

2.1 Support Building Extension: Captures the work associated with the extension of the support building, which was required to house the APOGEE and MARVELS instrument and to provide better storage of cryogens.

2.2 APOGEE Liquid Nitrogen Delivery System: Captures the work associated with building the system to transport cryogens from the storage area to the APOGEE dewar.

2.3 Cartridge Development: Captures the work associated with making the BOSS.

2.4 Cartridge Handling Development: Captures the work associated with the redesign of the equipment to manipulate the cartridges.

2.5 Cartridge and Plate Storage: Captures the work associated with the redesign of the plug room to store the cartridges.

2.6 Plate Marking Hardware Upgrade: Captures the work associated with marking the plates for efficient plugging in light of the multi-instrument plates.

2.7 Fiber Mapping Upgrade: Captures the work associated with building an adapter to allow the APOGEE and MARVELS fibers to interface to the standard fiber mapper.

2.8 Fiber Connector Development: Captures the work associated with developing the APOGEE and MARVELS fiber connectors.

2.9 Fiber Development: Captures the work associated with making the new science and guide fibers and anchor blocks.

2.10 Guider Development: Captures the hardware and software work associated with the replacement of the guide camera.

2.11 Operations Software Development: Captures the work associated with the redesign of the operations software, including the observing stations and the telescope control.

2.12 Requirements, Design, and Acceptance Reviews: Captures the work associated with the review life cycle of the infrastructure projects.

2.13 Documentation: Captures the work associated with documentation of the infrastructure projects.

2.14 Infrastructure Project Management: Captures the work associated with managing of the infrastructure projects.

3.3. SEGUE-2

3.1 Survey Strategy and Target Selection Development: Captures the work associated with design and implementation of the survey strategy and target selection for SEGUE-2.

3.2 On-mountain Quality Assurance Tools: Captures the work associated with the modification of the existing SEGUE QA tools to accommodate the new target selection and strategy.

3.3 Data Reduction Pipeline Development: Captures the work associated with modifications of the SEGUE pipelines to accommodate the new target selection.

3.4 Pipeline Quality Assurance Tools: Captures the work associated with development of QA tools for the SEGUE-2 dataset.

3.5 Data Reduction Pipeline Operations: Captures the work associated with reducing the SEGUE-2 data and validating the pipeline output.

3.6 Design and Acceptance Reviews: Captures the work associated with the review life cycle of the SEGUE-2 projects.

3.7 Documentation: Captures the work associated with the documentation of the SEGUE-2 pipelines, products, and data releases.

3.8 SEGUE-2 Project Management: Captures the work associated with management of SEGUE-2, including the PI and Survey Scientist.

3.4. MARVELS

4.1 Instrument Maintenance: Captures the work associated with maintaining the ET1 instrument, including support of the instrument scientist.

4.2 Survey Design and Strategy Development: Captures the work associated with the design and implementation of the survey strategy. As MARVELS is a monitoring project, the implementation of strategy continues into the operations phase, as the team must track the cadence of observations on each plate.

4.3 Target Selection Software Development: Captures the work associated with the design and implementation of target selection.

4.4 On-mountain Quality Assurance Tools: Captures the work associated with development of on-mountain QA tools to validate data quality.

4.5 Data Reduction Pipeline Development: Captures the work associated with development of the data reduction pipeline, including diagnosis of areas where performance is being impacted and implementation of improvements.

4.6 Pipeline Quality Assurance Tools: Captures the work associated with development of QA tools for the MARVELS pipeline by which the team can monitor survey quality and flag poor data.

4.7 Data Reduction Pipeline Operations: Captures the work associated with reducing the data and validating the pipeline outputs.

4.8 Design and Acceptance Reviews: Captures the work associated with life cycle reviews of the MARVELS project.

4.9 Documentation: Captures the work associated with the documentation of the MARVELS pipelines, products, and data releases.

4.10 MARVELS Project Management: Captures the work associated with management of MARVELS, including oversight by the PI and Survey Scientist.

3.5. BOSS

5.1 Spectrograph Upgrade: Captures the work associated with upgrade of the SDSS spectrograph, including the new detectors, dewars, grisms, optics, and dichroics.

5.2 Data Acquisition System Upgrade: Captures the work associated with upgrading the existing data acquisition system to work with the BOSS spectrograph.

5.3 BOSS Fiber Mapping Software Upgrade: Captures the work associated with upgrade of the fiber mapper to handle the new fiber system.

5.4 Survey Strategy and Target Selection Software Development: Captures the work associated with the design and implementation of the survey strategy and target selection software.

5.5 On-mountain Quality Assurance Tools: Captures the work associated with the upgrade of the on-mountain QA tools for the new spectrograph and target selection.

5.6 Data Reduction Pipeline Development: Captures the work associated with the upgrade of the data reduction pipeline.

5.6.1 Spectro-2D Software Development: Captures the work associated with the upgrade of the spectral extraction pipeline.

5.6.2 Spectro-1D Software Development: Captures the work associated with the upgrade of the redshift pipeline and other spectral analyses.

5.7 Pipeline Quality Assurance Tools: Captures the work associated with the development of tools to assess the data quality and flag bad data.

5.8 Commissioning and Acceptance: Captures the work associated with the commissioning and acceptance testing of the BOSS hardware and software systems.

5.9 Data Reduction Pipeline Operations: Captures the work associated with reduction and validation of BOSS data and pipeline outputs.

5.10 Design and Acceptance Reviews: Captures the work associated with life-cycle reviews of the BOSS project.

5.11 Documentation: Captures the work associated with documentation of the BOSS pipelines, products, and data releases.

5.12 BOSS Project Management: Captures the work associated with management of MARVELS, including oversight by the PI and Survey Scientist.

3.6. APOGEE

6.1 Instrument Design: Captures the work associated with the design of the APOGEE instrument.

6.2 Instrument Construction: Captures the work associated with the construction of the APOGEE instrument.

6.2.1 Optics: Captures the work associated with the construction of the camera, collimator, fold-flat, and fiber systems.

6.2.2 Electronics: Captures the work associated with the detectors and electronics.

6.2.3 Mechanical: Captures the work associated with the mechanical mounting of all elements and the construction of the cryostat.

6.2.4 Site Preparations: Captures the work associated with the preparation of the site to receive the instrument.

6.2.5 Integration & Testing: Captures the work associated with integrating the hardware components and doing all laboratory tests.

6.2.6 Instrument Documentation: Captures the work associated with documenting the instrument for facility status at APO and for public documentation of the survey.

6.3 Operations Software Development: Captures the work associated with instrument software and observing control software, including interfaces to the operations software hub.

6.4 Survey Strategy and Target Selection Software Development: Captures the work associated with the design and implementation of the survey strategy and target selection software.

6.5 On-mountain Quality Assurance Tools: Captures the work associated with development of tools for on-mountain assessment of data quality.

6.6 Data Reduction Pipeline Development: Captures the work associated with data reduction pipeline, including spectral extraction, calibration, radial velocity determination, and abundance determinations.

6.7 Pipeline Quality Assurance Tools: Captures the work associated with development of tools to validate the pipeline outputs.

6.8 Commissioning and Acceptance: Captures the work associated with commissioning the hardware and software and doing all acceptance tests.

6.9 Data Reduction Pipeline Operations: Captures the work associated with reduction and validation of APOGEE data and pipeline outputs.

6.10 Design and Acceptance Reviews: Captures the work associated with life-cycle reviews of the APOGEE project.

6.11 Documentation: Captures the work associated with documentation of the APOGEE survey and data products.

6.12 APOGEE Project Management: Captures the work associated with the management of the APOGEE project, including oversight from the PI and Survey Scientist.

3.7. Survey Observing

7.1 Facilities Operations: Captures the work associated with operations at Apache Point Observatory, including network systems, monitoring of facilities, site supplies and services.

7.2 Facilities Maintenance: Captures the work associated with maintenance of the observing systems at Apache Point Observatory and site maintenance.

7.3 SEGUE-2 Observing: Captures the work associated with night staff, instrument calibrations, and plate plugging for SEGUE-2.

7.4 MARVELS Observing: Captures the work associated with night staff, instrument calibrations, and plate plugging for MARVELS.

7.5 APOGEE/MARVELS Observing: Captures the work associated with night staff, instrument calibrations, and plate plugging for joint operations of APOGEE and MARVELS.

7.6 BOSS Imaging Observing: Captures the work associated with night staff and instrument calibrations for the BOSS imaging.

7.7 BOSS Spectroscopy Observing: Captures the work associated with night staff, instrument calibrations, and plate plugging for BOSS spectroscopy.

7.8 Apache Point Observatory Management: Captures the work associated with management of APO, including oversight from the site director and site manager.

7.9 Plate Design: Captures the work associated with design of plug plates from target lists.

7.10 Plate Drilling: Captures the work associated with the drilling of plug plates and shipping to APO.

3.8. Data Distribution

8.1 Computer Servers: Captures the work associated with the purchase and implementation of the large computer servers.

8.1.1 Science Archive Server: Captures the work associated with the purchase and implementation of the Science Archive Server.

8.1.2 Science Archive Mirror: Captures the work associated with the purchase and implementation of the mirror of the Science Archive Server.

8.1.3 Catalog Archive Servers: Captures the work associated with the purchase and implementation of the Catalog Archive Servers.

8.1.4 Catalog Archive Mirrors: Captures the work associated with the purchase and implementation of the mirrors of the Catalog Archive Servers.

8.2 System Administration: Captures the work associated with computer system administration on the archive servers.

8.3 Data Archiving: Captures the work associated with archiving of the raw data, including transfer to the Science Archive Server and Mirror, as well as an off-line backup at APO and tape backups at NERSC.

8.4 Data Assembly & Coordination: Captures the work associated with the assembly of the pipeline data products and target selection inputs into cohesive science-ready packages.

8.5 CAS Software Development: Captures the work associated with development of the Catalog Archive Server databases, including the development of the data schema, the data loaders, and the CAS interface.

8.6 CAS Adminstration: Captures the work associated with the administration of the database.

8.7 Quality Assurance Tools: Captures the work associated with development of tools to validate the data assembly and the CAS.

8.8 Data Releases: Captures the work associated with loading the data releases into the CAS and validating the results.

8.9 Design and Acceptance Reviews: Captures the work associated with life-cycle reviews of the data distribution project.

8.10 Documentation: Captures the work associated with documentation of the data distribution system and the data releases.

8.11 Help Desk: Captures the work associated with on-going expert support of the public data releases.

8.12 Community Training: Captures the work associated with programs to train the astronomical community in the use of the SDSS-III databases and data products.

8.13 Data Project Management: Captures the work associated with the management of the data distribution project, including the oversight from the data coordinator, the CAS lead, and the data project manager.

4. Development Projects: Control Procedures

The SDSS-III observing and data processing systems involve several new development projects, both in instrumentation and software. The Director has established an Instrument Review Board and a Software Review Board to oversee these projects. These are chaired by the Program Manager with membership comprised of experts involved in the SDSS-III development. Progress on each development effort will be reviewed by the Boards to track technical, schedule, and budget performance, enabling the Central Project Office to make any necessary adjustments in a timely manner. Key reviews will include external members. Release of funds for construction activities will be contingent on successful reviews at each stage. Documentation describing the newly developed systems from preliminary design to as-built components will be delivered to the project and baselined. The final products and documentation will be accepted by the Central Project Office when successfully measured against the baseline survey requirements.

Hardware projects will undergo a requirements, preliminary design, critical design, construction, commissioning, and acceptance life cycle under the direction of the respective team PIs, with programmatic accountability to the Central Project Office. The Instrument Board will recommend relevant hardware, interface control, acceptance, and documentation criteria and standards that will be employed by the developers and audited by the Central Project Office during the hardware development life cycles.

New (or modified) software required for SDSS-III will be developed to accepted standards and systems engineering practices recommended by the Software Board. We use strict source code version control, package management, and version tagging. Data models are clearly documented with change controlled by the Data Coordinator. SDSS-III data reduction pipelines will largely be written in IDL, as we have extensive infrastructure in place. The development lifecycles of the various SDSS-III software projects will be tracked by the Central Project Office for schedule, budget, and requirements adherence. In addition to unit testing done at the developer and team level, project-level reviews of software development projects will be undertaken at key design and development milestones, and at acceptance. The Software Board will organize and conduct these reviews, with external reviewers participating. Software will be developed using standards and systems engineering practices recommended by the Software Board.

Detailed development project schedules are being developed for each new instrument and its software. These schedules provide an important tool for the management of these activities and will be indexed to the baselined budgets and Work Breakdown Structure, so that schedule and budget are monitored in an integrated fashion. Progress in the largest instrument development effort, the APOGEE spectrograph, will be measured using earned-value management techniques. Beginning at the time of the CDR, the PI and Project Manager will provide the Program Manager an updated critical path analysis and report earned value and costs on a monthly basis during hardware and software development phases, in order to maintain centralized control and identify any developing delays early enough to take corrective action. This process was successfully used in the development of the MARVELS ET1 instrument. The BOSS spectrograph upgrades are smaller in scope and do not require this level of management oversight.

5. BOSS Development

BOSS is a spectroscopic survey that will target objects selected from the SDSS imaging catalog, including an additional 2000 square degrees to be completed during the first two years of SDSS-III. The existing SDSS spectrographs will be upgraded for BOSS, both to improve the throughput and to increase the number of fibers. The software pipeline will be revised and updated to accommodate the hardware changes. Hardware upgrades are managed by the Instrument Scientist, Natalie Roe, and software upgrades are managed by the BOSS PI, David Schlegel. There are weekly telecons to report progress, with minutes posted to the BOSS email archive. The BOSS wiki is used to document all technical designs, procurements and test results, with final drawings additionally stored on the APO server. Final designs are approved by the Instrument Scientist before fabrication, usually in consultation with the PI and others when appropriate. Many of the scientists and engineers involved in the original SDSS spectrograph design and fabrication are participants in the BOSS upgrade, providing the project with an extensive base of knowledge and experience.

5.1. Hardware Development

The BOSS spectrographs are an upgrade of the existing SDSS spectrographs, which are fed from fibers positioned using pre-drilled plates. Up to eight plates can be observed per night using the eight "dark time" fiber cartridges. For BOSS, the number of fibers per cartridge will be increased from 640 to 1000 in order to observe more objects per plate. In addition, the total throughput will be improved in both the blue (for the QSO Lyman-alpha forest studies) and the red (for luminous red galaxies, LRGs, up to a limiting redshift of z=0.7). Increasing both the number of objects per plate and the throughput is necessary for BOSS to meet its scientific goals of observing 1.3 million LRGs and the Lyman-alpha forest from 160,000 quasars. The implementation of the hardware upgrades will include eight new spectroscopic cartridges, four new cameras, and replacement of the dispersive elements with volume phase holographic (VPH) gratings, as well as several other elements in the optical chain. The eight "dark time" spectroscopic cartridges for BOSS will be outfitted with 1000 2-arcsec diameter science fibers. In addition, these cartridges will use 14 coherent fiber bundles for guiding with 7 arcsec diameters for guiding, and one fiber bundle with 25 arcsec diameter for field acquisition. The prototyping work for the science fibers was funded by the DOE and carried out by the BOSS team; however the final design and construction work falls under the Common Infrastructure upgrades (Section 9). The work on the cameras and optics is described here.

5.1.1. Cameras

Each of the four new cameras consists of a dewar, a single 4k x 4k, 15 um pixel CCD, and modified readout and control electronics and software. The custom dewars are designed and built by Michael Carr at Princeton, under the direct supervision of Jim Gunn, both of whom were responsible for the original SDSS spectrograph dewars. The final two corrector lenses in the optical train are mounted in the dewar with the final lens serving as the dewar window. These lenses will be specified and procured by Robert Barkhouser at JHU. The two blue cameras will have e2v Technologies CCD231-84 CCDs with enhanced blue sensitivity and expected read noise of 2 e-, while the two red cameras will house LBNL full-depletion CCDs with enhanced red sensitivity and expected read noise of 3 e-. The SDSS readout electronics will be modified for faster readout (due to increased pixel count) and to provide the specialized voltages required by the LBNL CCD. The electronics upgrade is designed by Jim Gunn (Princeton) and Connie Rockosi (UC Santa Cruz), both of whom were responsible for the original SDSS electronics, and the boards are fabricated and tested at Ohio State University under the supervision of Klaus Honscheid and at the University of Chicago Engineering Center.

The cameras underwent a preliminary design review on 22/23 January 2008 at LBNL. That review included all the principals for the dewars, CCDs, electronics, and optics. Technical progress is reported and discussed at the weekly BOSS Hardware Telecon. The mechanical drawings are posted in PDF format on the BOSS hardware wiki page, along with test results, procurement documents, and photos tracking progress.

The blue-side CCDs have been ordered from e2v Technologies, with delivery of three science-grade devices (including one spare) expected by May 2009. Recent communications from e2v suggest that an earlier delivery is possible. The red-side CCDs are fabricated at the LBNL MicroSystems Lab (MSL) and their industrial partner, DALSA Semiconductor. The first engineering grade LBNL 4k x 4k CCD has been delivered and tested (see below for more details), and a lot of 21 wafers intended for the BOSS devices has been delivered by DALSA. Fifteen wafers have been thinned, and are undergoing final processing in the LBNL MSL. The first five wafers have been completed and cold probed, and the results indicate that there are at least two promising science grade candidates.

A prototype dewar has been completed and passed the preliminary vacuum and thermal tests at Princeton in November 2008. This dewar was shipped to Apache Point in December 2008, where it was integrated with an engineering grade LBNL red CCD and the red-side modified electronics, including a new pre-amp board located in the dewar. Jim Gunn and Michael Carr (PU), Connie Rockosi (UCSC), and Natalie Roe, Armin Karcher and Bill Kolbe (LBNL) were present for the test, which was ultimately very successful after sorting out some of the usual prototype issues. The dewar performed as expected and the electronics delivered read noise below 3e- on all four readout amplifiers. Cosmic rays could clearly be seen in the images indicating good charge transfer efficiency. With this successful test completed, the “prototype” dewar will be used for one of the four cameras and three identical dewars will be fabricated; parts have already been ordered and machining is in progress. The prototype runs for the red-side electronics included enough boards to complete the red cameras, so they just need to be loaded and tested. The remaining electronics work for the blue side includes a pre-amp and one additional board, both of which are similar to those already completed for the red cameras and are expected to be complete in early 2009.

5.1.2. Optics

The optics upgrades include new dichroic beam splitters, new VPH gratings, new L7 and L8 lenses (mounted in the dewars), and re-spacing several of the existing optical elements in a new mechanical Central Optics Assembly (COA). The optical design work is being performed by Robert Barkhouser, and the opto-mechanical work by Stephen Smee, both at JHU. Barkhouser and Smee were involved in the original SDSS optical design.

The optical design underwent a preliminary design review on 18 April 2008 at JHU. Almost all of the action items identified at that review have been addressed, including a detailed study by Barkhouser of the ghosting from the VPH gratings. That study resulted in a re-design of the angles of the gratings and the prisms to send the light from the ghosts out of the active area. The optical studies are discussed at the weekly BOSS Hardware Telecon and in the BOSS mailing list. Final versions of these studies are posted on the BOSS hardware wiki page. The optical drawings are posted in PDF format on the wiki page, along with some photos tracking progress. Final designs are approved by the Instrument Scientist before fabrication, usually in consultation with the PI and others when appropriate.

The substrates for the two dichroic beam splitters have already been ordered and received. The coatings will be completed in the near future under a PO to JDSU that is currently in process at JHU, with delivery expected in April 2009. Each of the four dispersive elements consists of one VPH grating and two prisms. The substrates for the prisms have been ordered and received, and the sizing and polishing will be completed under a quote from Zygo that specifies a delivery of 11 weeks, putting delivery in mid-March. Prism coatings will follow and may be done at a different vendor. A prototype red VPH grating was ordered in October 2008 from Kaiser Optical Systems, Inc (KOSI) and we are still awaiting delivery. However, we are in regular contact with KOSI and understand the reasons for the delay. They are commissioning a new setup that is capable of exposing the large VPH gratings required by BOSS. Initial delays were due to the demands of other orders, and then various difficulties in the setup and commissioning of the new equipment contributed further delays. In December, KOSI made several exposures for the BOSS grating and we are hopeful that we can take delivery early in 2009 of both red gratings. We have also given them a PO for the delivery of the blue VPH gratings and they have quoted a delivery of 20 weeks, making this the critical path for the optical upgrade. The orders for the final corrector lenses will be placed in early 2009, and like the COA are not expected to be on the critical path.

5.2. Instrument Software

The instrument software component of the BOSS spectrograph upgrade entails replacing the instrument control microprocessor that runs the SDSS spectrographs with a more capable dual microprocessor system based on the one that runs the SDSS imaging camera. This dual system can drive the four-amplifier BOSS CCDs at the doubled readout speed, and adds increased flexibility that overcomes some of the operational limitations of the current SDSS spectrographs, most notably in observing efficiency by making it possible to send queries and commands to the instruments while they are reading out. The instrument control system upgrade has the following parts:

1) Modifications to the existing SDSS spectrographs microprocessor board so that it can function as the clock-generating microprocessor in the dual system.

2) Modifications to the existing SDSS camera "executive" microprocessor so that it can perform similar high-level functionality in the new BOSS dual-microprocessor system.

3) Modification of the existing clock generation microprocessor software and firmware to provide the 88 kHz readout for BOSS and the specific clocking requirements for the BOSS CCDs

4) Merging the more capable functionality of the SDSS imaging camera executive microprocessor software with the existing high-level SDSS spectrograph microprocessor software. This provides the capability to control multiple video processing and bias voltage generation circuit boards for each CCD, necessary to control the four-amplifier BOSS CCDs, with the specific housekeeping functionality of the spectrographs: liquid nitrogen fill, temperature monitoring, etc. Steps 1-4 are the minimum required for operating a four-amplifier BOSS CCD.

5) Upgrade of the software to convert to more BOSS-specific code in order to provide a user interface that is usable for engineering/commissioning, low-level debugging during operations, and that meets the interface requirements of the SDSS-III operations software. These upgrades provide all the functionality required for normal observing and performance monitoring.

The instrument software work is under the purview of the BOSS Instrument Scientist. The development work must be done at Apache Point, in coordination with the site staff. The final product must integrate with the SDSS-III operations software.

The coding will be done by Connie Rockosi, who did the coding for the SDSS spectrographs in the late 1990s. A successful pass through steps 1-4 was completed in time to run an engineering grade fully-depleted LBNL CCD for the December, 2008, dewar test at Apache Point. Step 5 will require adding BOSS-specific functionality and writing to the SDSS-III operations software interface standards currently under development.

5.3. Pipeline Software Development

The spectroscopic code consists of several pipelines that share libraries. The on-the-mountain reductions (Son-of-Spectro, or SoS) is a real-time reduction pipeline that gives the observers' feedback on when S/N is achieved on a given plate. The Spectro-2D pipeline is the more sophisticated version that reduces the 2-dimensional spectroscopic images to 1-dimensional spectra, errors, and other associated vectors. The Spectro-1D pipeline determines the classification and redshifts of objects based upon the 1-d spectra. There are currently no pipelines for generating fake data.

The details of the data will change between SDSS and BOSS with the upgrade from 640 fibers to 1000 smaller fibers, new fiber and fiber bundle spacings, the replacement of the CCDs with larger formats and different characteristics, and the replacement of several optics (notably the dichroics and gratings) which will change the wavelength mappings and 2-d PSFs. All three pipelines must be upgraded to accept the BOSS data, which includes quite a number of hard-wired features. A simple, simulation pipeline will generate fake data such that the basics of the pipeline upgrade can be verified before commissioning in summer 2009. In parallel with this effort, several important upgrades for the Spectro-2D and Spectro-1D pipelines will be developed in order to improve the LRG redshifts and to enable the full science return from the QSO Lyman-alpha forest.

5.3.1. Simulation Pipeline

A simulation pipeline for generating fake data will be written in the 2008-2009 academic year by Lauren Anderson at LBL. This will be under the direction of the PI. Relevant inputs for the fake LRG spectra and fake QSO spectra will be provided by the galaxy and QSO working groups, respectively. For the LRGs, this demands a realistic distribution of redshifts and magnitudes, and realistic stellar population models. For the QSOs, this demands a realistic distribution of redshifts and magnitudes, underlying QSO spectral templates from N. Suzuki's PCA decompositions of spectra, and J. Hennawi's models of Lyman-alpha forest.

5.3.2. Real-Time Reduction Pipeline

The minimal upgrade of the SoS, Spectro-2D and Spectro-1D pipelines will be made by Lauren Anderson in Jan-June 2009. This will be under the direction of the PI. These will include simulated flat-field images, arc exposures and science exposures. This code port will maintain the SDSS-I file format, and the log-wavelength mapping for the co-added spectra. The tools for accessing the outputs (READSPEC) will also be modified. This code will be tested with the fake data generated by the simulation pipeline.

The real-time SoS reduction code measures the S/N of objects at fiducial magnitudes at the end of each exposure. It also tracks the sky brightness and test for the recurrence of known problems with the spectrographs. The scientific decision is the determination of the new S/N thresholds for the BOSS targets. The blue S/N thresholds will be driven by the QSO science, and the red S/N thresholds by the galaxy science, but they will be jointly tuned to give comparable exposure times since targets for both programs are observed in parallel. These S/N decisions, and SoS code changes, will be made by a postdoc (tentatively at Princeton), under the direction of the PI.

5.3.3. Pipeline Upgrades

The expectation is that the full requirements of the demanding QSO Lyman-alpha forest analysis will not be met with the minimal modifications to the SDSS spectroscopic pipeline. In order to meet requirements, a number of modular upgrades to the pipeline are planned during survey operations.

The most significant module is a new 2-D extraction pipeline, which will be implemented using the Bolton & Schlegel algorithm. This will make use of the full 2-dimensional PSF and will track errors and covariances in the extracted data. The other modules are the pre-processing steps, pixel flats, PSF modeling, scattered light modeling, fiber flats, wavelength-calibration (from arc and sky lines), sky-subtraction, flux-calibration (including telluric correction). QA will be output appropriate for tracking down problems in the data and in the pipeline. This work will be performed by an LBNL postdoc, with likely input from interested parties at other institutions.

The Spectro-1D pipeline takes the 1-D spectra as input, classifies objects, fits redshifts, and measures a number of basic galaxy properties such as emission line strengths and velocity dispersion. This pipeline will be upgraded to take the new Spectro-2D data format as inputs, working off the 1-D spectra on the native wavelength mapping of the individual exposures. (The SDSS-I Lyman-alpha forest analysis was done by P. McDonald using the individual exposures, rather than the publicly-released co-added spectra.) The QSO redshift fits will be improved with a more proper treatment of the Lyman-alpha forest region and enforced dependence upon redward features when available. These upgrades will be performed by a postdoc, in coordination with the Lyman-alpha forest working group.

The pipeline upgrades will necessarily require a few, discrete changes to the BOSS spectroscopic data model. Those changes will be submitted for approval by the Data Coordinator.

5.3.4. Spectro Templates

The spectral templates must increase the wavelength range to that of the BOSS spectrographs, which extend coverage both further into the UV (3500 Ang) and into the near-IR (1.05 microns). Although the BOSS primary science is only LRGs and QSOs, the templates must include the full suite used in SDSS-I in order to avoid spectroscopic mis-classifications by the pipeline. The initial updates of these templates will use the BOSS commissioning data. As in SDSS-I, subsequent updates of these templates will be possible as more spectra are taken (especially of more rare objects) and any problems in the templates are uncovered. The initial templates will be constructed by an LBNL postdoc. Any subsequent updates will be made by the graduate students tracking the LRG and QSO success throughout the survey.

5.3.5. Automated Reductions and Data Packaging

The Spectro-2D and Spectro-1D reductions are largely an automated process. It will be overseen by an LBNL postdoc with responsibility for ensuring all of the data are sequenced properly through the pipelines. On a daily basis, the copying of the data will be verified by the SDSS-III Data Group (currently Ben Weaver). The LBNL postdoc will ensure that bad exposures are marked as such, and will read the observer logs for problems. This person will track down problems reported in the night logs, by the real-time reductions, and by the full reductions. Problems not related to weather will be tracked through the SDSS-III ticket system.

The spectroscopic data will be packaged in the format necessary for the Catalog Archive Server, including any auxiliary pipeline outputs. Known extant problems will be characterized and documented for the formal data releases. This is a 25% effort to be done by a postdoc (tentatively at Princeton).

5.4. Target Selection Implementation

The BOSS target selection will be made from SDSS photometry. This includes the LRG targets, QSO targets, and the calibration fibers (F stars for spectro-photometry and blank skies). The baseline target selection is included in the BOSS requirements document. Amendments will be made through the Change Control Board, with final changes expected during BOSS commissioning.

The BOSS LRG targets are very well-understood from the SDSS and 2SLAQ surveys. There is an opportunity to tune the details of the target selection to allow broader studies of galaxy evolution in addition to the core cosmology science program. The galaxy working group has convened a meeting in January 2009 (with 25 participants) to consider any amendments to the baseline LRG targets. If the commissioning phase demonstrates any problems in achieving the required completeness, the target selection will be amended to avoid problematic regions in color-magnitude space. No amendments are expected to be made to the LRG targeting after the commissioning phase, as that would unnecessarily complicate the BOSS primary science analysis.

The BOSS QSO targets are less-well understood to the limiting magnitude of i=22, and are expected to suffer 60% contamination from stars and low-z QSOs. The QSO working group has embarked on a thorough program to construct a "truth table" of QSO targets on several footprints across the sky. This includes additional MMT spectroscopy in October 2008 and January 2009, led by Adam Myers and Gordon Richards. This truth table will be used to test the QSO selection algorithms developed by Gordon Richards, Joe Hennawi, and Robert da Silva. Based upon these data, a recommendation will be made from the QSO working group to the PI and the Change Control Board to amend the QSO target selection.

The BOSS QSO primary science (Lyman-alpha forest analysis) uses the QSOs as a back-light to cosmic structure, and does not require uniform selection throughout the survey area. This will allow the flexibility to amend the QSO targeting, should future data allow more efficient rejection of contaminating objects. Possible future data sets include improved magnitudes from Pan-STARRS, variability information from Pan-STARRS, infrared magnitudes from UKIDDS, and a recently-approved program to extend the FIRST radio catalog to the BOSS imaging in the Fall sky. A meeting of the QSO Lyman-alpha forest working group at Ohio State in March 2009 will recommend any changes to the baseline targeting algorithm on core scientific grounds. If there are significant improvements to be made to the QSO targeting, changes will be requested by the QSO working group to the PI and the Change Control Board in time for each successive year's observing (before summer shut-down).

The LRG target selection code has been written by Nikhil Padmanabhan. The QSO target selection code has been written by Robert da Silva. Updates to these codes will be made by members of the respective working groups, under the direction of the PI.

The success of the LRG and QSO target selection algorithms over the course of the survey will be monitored by graduate students at Berkeley, under the direction of the PI. The LRGs must conform to the redshift distribution defined in the requirements document, and achieve the required redshift success rate. The QSOs must meet the required number density on the sky, achieve the required redshift success rate, and meet the S/N requirements as a function of magnitude.

5.5. Review & Schedule Milestones

22/23 January 2008	BOSS Camera Upgrade Preliminary Design Review
18 April 2008	BOSS Optics Upgrade Preliminary Design Review
February 2009	BOSS Camera and Optics Upgrade Critical Design Review
March 2009	BOSS Target Selection and Survey Strategy Review
May 2009	BOSS Data Reduction Preliminary Design Review

Additional milestones and possible reviews for the data reduction pipeline will be considered at the May 2009 design review.

5.6. Commissioning Plan

The dewars, CCDs and electronics boards will be delivered to APO and integrated there. Each CCD will be installed in its dewar and independently tested with its readout electronics, using the same test stand employed for the prototype camera test. After a successful test, the cameras will be ready for a drop-in replacement into the existing SDSS spectrographs.

The optical elements, including the dichroics, gratings and new central optical assembly in which they will be mounted, will be delivered to JHU for integration and testing. Then the individual elements will be shipped to APO where they will be reassembled, and again it will be a drop-in replacement for the existing spectrograph optics.

Commissioning the spectrographs requires the repeat of a number of alignment tests and taking of calibration data performed for SDSS-I in Fall 2000. The Son-of-Spectro tools, which did not exist in 2000, will make these tasks more straightforward in 2009. The alignment tests include calibration of the collimator positions for proper wavelength and fiber coverage and focusing of the spectrographs. The calibration data include the construction of stacked bias exposures, the construction of stacked dark exposures, and characterization of the scattered light on the CCDs (using sparsely-plugged plates). Testing that throughput requirements are met will be done with the SPTHROUGHPUT procedure on the first plate observed in good-seeing, photometric conditions. Testing the instrument flexure in each of the 4 cameras will be done with a flat-field and arc exposures taken at a series of telescope altitude and azimuth positions. These tests will be performed by Lauren Anderson and any BOSS software postdocs under the direction of the PI and Instrument Scientist.

5.7. Acceptance Test

5.7.1. Instrument Acceptance Tests

Instrument acceptance tests will be undertaken under the management of the Instrument Scientist. Several of the hardware leads will be present or available: R. Barkhouser or S. Smee (optics), C. Rockosi or J. Gunn (electronics), K. Honscheid (DAQ). The calibration and on-sky test data will be taken by the Apache Point observing staff. The PI will be present with several members of the BOSS team to verify that each of the hardware requirements in the requirements document are met. The Son-of-Spectro pipeline tools will be used for these measurements, which will require the presence or availability of software developers to make any last-minute code changes that should be necessary.

5.7.2. Software Pipeline Acceptance Tests

The spectroscopic software pipeline acceptance tests will be undertaken under the management of the PI. The initial delivery of the pipeline will be required to reduce existing SDSS-I data and mock BOSS data, returning either the correct redshifts with normally-distributed errors, or assigning redshifts as "unknown" for objects where the S/N is inadequate. It is not expected that the initial version will satisfy the Lyman-alpha forest noise requirements from the mock data.

The BOSS commissioning data will cover regions of the sky with known LRG redshifts from the 2SLAQ survey and known QSO redshifts from MMT data currently being taken for that purpose. The pipeline will be required to recover the redshifts for these two categories of objects with normally-distributed errors. The failure rate must not exceed that specified in the requirements document. Later versions of the BOSS pipeline will include upgrades or replacement of discrete components. The pipeline developers must demonstrate that each new version meets the redshift accuracy requirements, and exceeds the previous version in the Lyman-alpha forest noise tests.

Code versions for use in the formal data releases will be certified as such by the Survey Scientist and PI after being run on all the data included in those and previous data releases. In addition, the developers will identify and attempt to understand objects whose recovered redshifts significantly change between major code releases. Tools exist from SDSS-I to make such comparisons straight-forward. The collaboration members will be encouraged to pursue independent tests of the reductions and identify specific problems, which would be submitted through the ticketing system. The data reduction pipelines in SDSS-I and SDSS-II benefited greatly from such collaboration input.

6. APOGEE Development

6.1. Hardware Development

The APOGEE instrument team will design, build and commission the APOGEE spectrograph on the Sloan 2.5-m telescope. This instrument will be a cryogenic, near-infrared, fiber-fed spectrograph based upon a volume-phase holographic (VPH) dispersive element that provides spectral coverage from 1.51 to 1.68 um at a mean resolution of 24,000 for 300 fiber inputs arranged along a pseudo-slit. The spectrograph will conform technically to an opto-mechanical design capable of delivering spectra of stars in the numbers, throughput (efficiency) per spectrum, resolution, wavelength coverage and sampling specified in the APOGEE Science Requirements Document (SRD), with a goal of achieving routine survey operations starting in 2011 Q2.

6.1.1. Instrument Development Management

The management of the development of the APOGEE instrument is the responsibility of the APOGEE Instrument Scientist (Michael Skrutskie). The Instrument Scientist will develop a detailed plan for the design, fabrication, testing and validation, and delivery of the instrument to the telescope. The instrument scientist will be assisted by a Deputy Instrument Scientist (John Wilson) in all of these efforts. The Instrument Scientist will work with the APOGEE Project Manager to establish a schedule of reviews and milestones consistent with an on-time delivery of the system in order to support APOGEE operations.

6.1.2. Instrument Performance Requirements

APOGEE will be designed and constructed to meet the requirements specified in the APOGEE SRD. The instrument will undergo multiple reviews throughout its development cycle, including Conceptual Design Reviews, Preliminary Design Reviews, Critical Design Reviews, Pre-ship Reviews, and an Operation Readiness Review. The APOGEE Instrument Scientist has responsibility for directing the testing and evaluation of the APOGEE instrument and reporting those results to the APOGEE Principal Investigator (Steven Majewski). The Principal Investigator has responsibility for evaluating the instrument performance in the context of the SRD requirements and certifying that the instrument is ready for operations.

6.1.3. Instrument Design

Integrated Instrument Design

The APOGEE Instrument Scientist has responsibility for the integrated APOGEE instrument design. That design will consist of modular components, the design for some of these components will be procured from outside the APOGEE instrument team. The APOGEE Instrument Scientist will coordinate the efforts of the internal instrument team and external vendors to assure that the overall instrument design is feasible and will meet the scientific requirements of the APOGEE project.

Optical and Opto-mechanical Design

The APOGEE Instrument Scientist and Deputy Instrument Scientist will have responsibility for establishing an optical design for the APOGEE instrument. Major portions of this optical design, for example the large refractive camera optics, will be contracted to outside vendors. For this contract work the APOGEE instrument team will have responsibility for setting specifications for the performance of the optical systems and for communication with vendors during the development effort. The APOGEE instrument team, in conjunction with the APOGEE Project Manager, will work with vendors to establish milestones and delivery schedules consistent with the integrated project schedule.

A breakdown of remaining tasks is:

- Perform optics and opto-mechanical design of a custom refractive camera to image the dispersed spectrum of APOGEE onto the science detectors. Perform tolerancing and ghost analysis. This work will be conducted by New England Optical Systems (NEOS; Marlborough, MA,) via contract.
- Provide thermal analysis and consulting to NEOS to ensure opto-mechanical design for refractive camera performs as required after thermal cycling between room temperature and 77K. This work will be performed by Tom O'Brien (OSU) and Steve Smee (JHU) under contract.
- Provide stray light analysis of optical design to assist in identification of unforeseen sources of stray light (e.g., lens/lens barrel interfaces, zero'th order through the VPH, etc.). This work will likely be performed by Breault Research (Tucson, AZ).
- Design and specify requirements for the custom VPH Grating. This work will be done by Robert Barkhouser (JHU), under contract, and John Wilson and Michael Skrutskie.
- Design the VPH opto-mechanics mount ensuring that the optic cools to cryogenic temperatures safely and that its temperature profile is within specifications. Work to be done by a contractor.
- Perform the optical design of the fore-optics, to include the fiber pseudo-slit, collimator and fold-mirror. This work is to be done by Robert Barkhouser (JHU), under contract, and John Wilson.
- Perform opto-mechanical design of fore-optics, including fiber pseudo-slit, collimator and fold-mirror. The collimator shall include a mechanism to adjust the tilt to allow small dithering of the spectrum (in the dispersion direction) for optimal sampling and tip in the spatial direction (for flat field calibration). Work to be done by a contractor.

Cryostat Design

The APOGEE Instrument Scientist and Deputy Instrument Scientist have responsibility for overseeing development of a cryogenic enclosure for the APOGEE instrument. The enclosure conceptual design will be executed under contract to an independent mechanical design contractor. The APOGEE instrument team will have responsibility for providing specifications for the cryostat based on the opto-mechanical requirements for the APOGEE project and for communicating with the vendor during the development of the design. The instrument team will have responsibility for assuring that the delivered design will meet the requirements for housing the APOGEE instrument. Upon establishing a valid conceptual design the instrument team will seek vendors who will execute a final design and deliver a finished cryostat to the APOGEE project.

A breakdown of remaining tasks is:

- Perform the rough mechanical design of a custom cryostat to house the APOGEE instrument. The design is to be at a level suitable for writing specifications in anticipation of a commercial procurement for the cryostat fabrication. Work to be done by Chuck Henderson (Independent Contractor) and Basil Blank (PulseRay Machining and Design).

Fiber System Design

The APOGEE Instrument Scientist and Deputy Instrument Scientist will be responsible for delivering APOGEE instrument fiber requirements to the SDSS-III infrastructure team and working with the infrastructure team to produce fiber harnesses, fiber couplers, and the main fiber run to the APOGEE enclosure. The infrastructure team has primary responsibility for the design and fabrication of all fiber system elements from the telescope focal plane to the fiber entrance into the APOGEE laboratory in the support building. The APOGEE instrument team has lead responsibility, with assistance from the infrastructure team where possible, for the vacuum interface of the fibers into the APOGEE cryostat and development of the cryogenic fiber pseudo-slit.

A breakdown of remaining tasks is:

- Perform the opto-mechanical design of the fiber system from the custom Gang-connector at the Sloan Telescope to the fiber pseudo-slit. Work to be done by the instrument team with French Leger (APO).
- Identify suitable vendors for fiber assembly fabrication (Michael Skrutskie and John Wilson).
- Design custom vacuum feedthroughs for the fiber assemblies to penetrate into the cryostat. Feedthrough design must minimize stress on fibers to mitigate focal ratio degradation (FRD). Work to be done by contractors and the instrument team.
- Design custom V-groove blocks for the fiber pseudo-slit so fiber launch at the correct optical orientation. Minimize stress on fibers to mitigate FRD. Work to be done by contractors and the instrument team.

Electronics and Array Mount Design

The APOGEE Instrument Scientist and Deputy Instrument Scientist have responsibility for overseeing the design of the APOGEE readout electronics system and the physical mounting of the arrays within the APOGEE cryostat. The readout electronics architecture will be established and implemented by the APOGEE instrument team. The array mount design and fabrication will be contracted to Erick Young (UA), with the instrument team having responsibility for establishing the performance requirements for the mounting system.

A breakdown of remaining tasks is:

- Specify and procure read-out electronics for the three Teledyne H2RG detectors. This work will be performed by Matt Nelson. Expect to baseline Leach electronics.
- Design cold fan-out board system. This work will be performed by Matt Nelson. Work with Erick Young (UA) regarding cabling plan.
- Design cabling system from fan-out board to warm hermetic connectors. This work will be performed by Matt Nelson, John Wilson and the instrument post-doc.
- Design a mosaic mount for three Teledyne H2RG detectors. Each detector shall possess independent (manual) piston and tip-tilt degrees of freedom. Work with NEOS so mount attaches to the back of the refractive camera. This work will be performed by Erick Young (UA), under contract.

6.1.4. Component Procurement

The APOGEE Instrument Scientist and Deputy Instrument Scientist will have responsibility of overseeing all aspects of procurement of the primary components of the APOGEE instrument system. In conjunction with the APOGEE Project Manager the instrument team will interact with

ARC and/or UVa procurement to assure that the required system elements are delivered on a schedule consistent with the integrated project schedule.

A breakdown of remaining tasks is (with all procurements the responsibility of the Instrument Scientist and Deputy Instrument Scientist working with the APOGEE Project Manager):

- Procure the custom refractive camera. The vendor will work with the APOGEE team to determine whether the VPH can be attached directly to the front of the camera and the detector mosaic assembly can attach directly to the rear.
- The instrument team will work with Tom O'Brien (OSU) and Steve Smee (JHU), who will provide thermal analysis and consulting on opto-mechanical designs to ensure the optics stay within tolerance despite thermal cycling between room temperature and 77K. Prior to acceptance and final milestone payment, the vendor, with assistance from the APOGEE team and using a cryostat provided by the APOGEE team, will cycle the camera to cryogenic temperatures to ensure acceptable performance. This work is expected to be performed by New England Optical Systems (NEOS; Marlborough, MA,) via contract.
- Fabricate and test the custom VPH Grating. This work is expected to be procured through competitive bid between Kaiser Optics and Wasatch Photonics.
- Collimator fabrication is expected to be competitively bid. Multiple firms are capable of this work.
- Fabrication of the fold mirror is expected to be competitively bid. Multiple firms are capable of this work.
- Design (to shop drawing level), fabricate and vacuum test custom cryostat. This work is expected to be performed by PulseRay Machining and Design.
- Procure raw fiber for APOGEE. This is expected to be FIP120170190 from Polymicro.
- Procure fiber assembly for APOGEE. This is expected to be competitively bid. Multiple firms are capable of this work although past Sloan Survey experience has shown that C-Technologies provides high quality assemblies at a competitive price.
- Fabricate detector mosaic mount assembly. This work will be performed by Erick Young (UA), under contract.
- Procure Leach Electronics for array read-out electronics.
- Procure custom fan-out board.
- Procure custom cryogenic cable assemblies and hermetic feedthroughs.

6.1.5. System Integration

The APOGEE Instrument Scientist and Deputy Instrument Scientist have responsibility for directing the assembly of the procured and fabricated components of the APOGEE instrument into a complete and functional instrument ready for delivery to the telescope. Instrument assembly will occur at UVA.

6.1.6. System Documentation

The APOGEE Instrument Scientist and Deputy Instrument Scientist have responsibility for documenting the instrument to allow full operation, trouble-shooting and upkeep while the survey is in operation. John Wilson and the instrument post-doc will lead documentation development.

6.1.7. Performance Validation

The APOGEE spectrograph must perform in a manner consistent with the science requirements for the project. The APOGEE Instrument Scientist and Deputy Instrument Scientist will oversee the validation of both the individual components of the instrument prior to integration as well as the integrated instrument.

Component Level Validation

Individual components of the APOGEE spectrograph may be tested/validated in conjunction with their vendors or in-house within the APOGEE project. In either case the APOGEE Instrument Scientist and Deputy Instrument Scientist carry responsibility for establishing validation standards, overseeing testing, and evaluating the performance of the individual components of the APOGEE instrument.

A breakdown of remaining tasks is:

- VPH optical testing (at vendor) and acceptance testing by the APOGEE team (at UVA or JHU Instrument Development Group, under contract).
- Cryogenic testing of a test VPH (not full-scale) with tilted fringes will be conducted at JHU in collaboration with Kaiser Optics.
- Cryogenic testing of the full-scale VPH will be performed to prove survivability. This testing will occur in the instrument cryostat at UVA.
- Cryogenic testing of the full-scale VPH may be performed to prove optical efficiency. This testing will require a cryostat with sufficient size and optics windows for passing collimated metrology beams in/out of the cryostat. If necessary we will contract to NASA Goddard and/or a commercial company for this testing.
- Fused silica blank testing may be conducted to prove opto-mechanical radial/axial mounting techniques. If the testing is deemed necessary, it will be performed by NEOS.
- Both the fold and collimator will tested by the vendor that fabricates each optic. The assembled camera optical performance will be tested before and after the cryogenic crack test by NEOS. Test results showing that the optics/camera meet specifications will be provided to UVA.
- Fiber testing to check the FRD of 40-meter fiber runs at 1.6 micron and FRD of prototype vacuum feedthrough to inform optical and fiber design will be conducted. This work is ongoing at UVA, currently by Sophia Brunner, John Wilson and Jeff Crane (OCIW).
- Fiber testing of prototype V-groove block will also be conducted at UVA to ensure minimal addition to FRD.
- Fiber assembly testing (FRD, throughput, and feedthrough vacuum integrity) will be accomplished at UVA upon receipt of finished assemblies. Work to be done by the instrument post-doc.

- Cryostat vacuum and cryogenic testing will be conducted upon arrival at UVa. Work to be done by John Wilson, Matt Nelson and the instrument post-doc.
- Fore-optics testing will ensure fiber orientation is satisfactory. Work to be done by John Wilson and the instrument post-doc.

Laboratory End-to-end Validation

The APOGEE Instrument Scientist and Deputy Instrument Scientist carry responsibility for establishing validation standards, overseeing testing, and evaluating the lab performance of the integrated APOGEE instrument.

- After passing its lab validation tests, the PI will recommend to the Central Project Office (CPO) that the instrument be shipped to the mountain for commissioning.
- The CPO will either accept the recommendation or request further validation tests or improvements.

Telescope Performance Validation

- The instrument will pass the acceptance tests outlined in Section 6.7 below.
- Hardware and instrument software (see Section 6.2 below) documentation will be delivered in sufficient detail to allow the start of survey operations.
- When on-mountain acceptance tests are passed the PI will recommend acceptance to the CPO, who will either accept the recommendation or request further validation tests or improvements.

6.2. Instrument Software

APOGEE requires software interfaces to

a) clock and read out the three HAWAII-2RG arrays

b) format and store raw data

c) control moving mechanisms within the APOGEE cryostat

d) collect instrument housekeeping information

e) control all of the above functions through the SDSS instrument server software.

The APOGEE Instrument Scientist and Deputy Instrument Scientist have responsibility for developing the specifications and overseeing the development of the instrument software interfaces. The software development work will be carried out at the University of Virginia. The instrument output will be based on a standard FITS format, suitably structured for direct input into both the APOGEE quality assurance (AQuA) and APOGEE pipeline reduction software codes.

Production Tasks (all activities carried out at UVa under management of Instrument Scientist and Deputy Instrument Scientist)

- Develop end-to-end readout test configuration using HAWAII-2RG bare multiplexer integrated with readout electronics.
- Produce working array controller code and test with HAWAII-2RG bare multiplexer, ultimately using production fanout and test cabling.

- Optimize controller code and fanout/cabling configuration to optimize read noise behavior.
- Produce server software for SDSS-III environment/hub.
- Develop code to format output data in FITS format with header information consistent with the requirements of the APOGEE pipelines.
- Integrate opto-mechanical and housekeeping functionality into server.
- Incorporate APOGEE software on-site with SDSS-III hub.
- Validate performance with science-grade detector arrays installed within APOGEE instrument.
- Develop instrument software documentation to allow instrument operation, trouble-shooting and upkeep.

6.3. Pipeline Software Development

The APOGEE pipeline software development breaks down into two broad categories: basic spectral reduction and derivation of stellar parameters. The PI and Survey Scientist (Ricardo Schiavon) will appoint APOGEE Software Teams responsible for overseeing software development in both categories. Each Software Team will have a Team Leader, who will report to the PI and Survey Scientist. Each software pipeline category (i.e. Team) may, if deemed necessary in the future, have specific tasks that will each be assigned to additional Task Leaders responsible for their definition, oversight and timely execution.

The appointed Software Teams, working with the APOGEE PI and Survey Scientist, will oversee development of software that will execute the following general tasks required to convert the APOGEE instrument output images to final data products suitable for public data release: (1) The APOGEE Spectral Extraction and Calibration Software (ASECS) will patch the output from the three detectors into a single 2D image, optimally extract one-dimensional spectra from the latter, wavelength calibrate the resulting 1D spectra, correct them for telluric absorption and instrumental response, and combine them with previously archived one-dimensional spectra for the same targets; (2) the APOGEE Radial Velocity Pipeline (ARVP) will determine the radial velocities for each target (most likely via cross correlation with the spectra of radial velocity standards); and (3) using the one-dimensional APOGEE spectra and available photometry, the APOGEE Stellar Parameters and Chemical Abundance Pipeline (ASPCAP) will derive for each star the basic atmospheric parameters (log g, [Fe/H], T_{eff}) and the abundances of the multiple chemical elements outlined in the SRD**,** to the precision discussed in that document, commensurate with the net *S/N* of the spectra.

In addition, the Software Teams will collectively assemble the APOGEE Quality Assurance (AQuA) Software, which will be a “quick-look” version of the other software packages to be run at APO by the Observers in order to obtain a real time assessment of the instrument output.

There are two budgeted postdocs, PD1 and PD2, tasked primarily with writing APOGEE reduction pipeline software and its documentation who will work under the appropriate Team Leaders for any given task.

APOGEE’s ultimate pipeline data deliverables will be: (a) fully wavelength-calibrated, telluric-absorption corrected, one-dimensional spectra for each star targeted by APOGEE, taking into account all accumulated integrations for each star; (b) a single radial velocity for each star, with an associated uncertainty, (c) the basic atmospheric parameters (log g, [Fe/H], T_{eff}) for each star, with uncertainties, and (d) the abundances of the multiple chemical elements outlined in the APOGEE SRD with associated uncertainties, to the precision discussed in that document.

6.3.1. APOGEE Spectral Extraction and Calibration Software (ASECS)

The ASECS will be developed by APOGEE PD1 under the supervision of the ASECS Team Leader, and with the participation of the ASECS Software Team. A breakdown of these tasks, with the primary responsibilities, is as follows:

- Appointment of ASECS Software Team Leader (PI, after consultation with Survey Scientist).
- Adaptation of existing SDSS-I and -II codes for spectral extraction.
  - Determine whether the presently existing SDSS flat-field lamp has adequate flux in *H*-band to be useful for APOGEE (ASECS Team Leader, PD1).
  - Determination of APOGEE requirement for relative continuum/line fluxing (ASECS Team Leader).
  - Measurement of the scattered light properties of the APOGEE spectrograph (PD1, with assistance of instrument team).
  - Determination of the procedure to account for contamination of adjacent spectra by scattered light, perhaps adapting SDSS-I and -II procedures, if adequate (PD1, ASECS Software Team).
  - Software development (PD1, with contributions from members of ASECS Software Team).
- Adaptation of SDSS code for sky subtraction.
  - Determination of the best procedure to account for spatial variations of sky background (PD1, with ASECS Software Team).
  - Decision on whether different airglow lines vary in tandem or whether different transitions vary differently both spatially and as a function of time (PD1).
  - Determination of how many sky fibers are needed to sample adequately sky emission across FOV (PD1 undertakes tests, all Software Teams discuss, final requirement made by Survey Scientist).
  - Software development (PD1).
- Adaptation of SDSS code for telluric division.
  - Determination of ideal spectral type/color for telluric standards (ASECS Software Team).
  - Software development (PD1).
- Adaptation of SDSS code for wavelength calibration.
  - Determination of the number of telluric lines usable for accurate wavelength calibration (PD1).
  - Software development (PD1).
- Instrument response calibration.
  - Establish procedure for instrument spectral response measurement on the basis of commissioning data (ASECS Software Team). Software development (PD1) to generate code for instrumental response extraction using standard star data (likely using mostly available software; PD1).
  - Software development (PD1) to correct APOGEE spectra from instrumental response, as necessary.

    - If instrumental response correction is deemed unnecessary for routine pipeline correction, decide on plan for periodically monitoring the instrument health (absolute throughput and also the relative shape of the spectral response; ASECS Software Team). Implementation by PD1 in cooperation with instrument team.
- Documentation of ASECS software (PD1, proofread and edited by ASECS Team Leader, approved by Survey Scientist).

6.3.2. APOGEE Radial Velocity Pipeline (ARVP)

The ARVP will be developed by APOGEE PD1 under the supervision of the ARVP Team Leader, and with the participation of the ARVP Software Team. A breakdown of these tasks, with the primary responsibilities, is as follows:

- Appointment of ARVP Team Leader (PI, after consultation with Survey Scientist).
- Specification of a set of RV standards with suitably wide range of spectral properties that will be used for RV observations and that will provide cross-correlation templates. The periodicity of the RV standard observations will depend on the spectrograph wavelength solution stability, which will be established during commissioning (ARVP Software Team, PD1).
- Development of cross-correlation software for measuring RVs and determining the precision as a function of S/N and spectral type (PD1).
- Documentation of ARVP software (PD1, proofread and edited by ARVP Team leader, approved by Survey Scientist).

6.3.3. APOGEE Stellar Parameters and Chemical Abundance Pipeline (ASPCAP)

The APSCAP will be developed by APOGEE PD2 under the supervision of the APSCAP Team Leader, and with the participation of the ASPCAP Software Team. It is anticipated that this software area will require several distinct Task Leaders. A breakdown of these tasks, with the primary responsibilities, is as follows:

- Appointment of ASPCAP Team Leader (PI, after consultation with Survey Scientist).
- Appointment of separate ASPCAP Task Leaders for (1) spectral synthesis, (2) stellar parameters, and (3) chemical abundances (Survey Scientist, after consultation with the ASPCAP Team Leader).
- Spectral synthesis and line list preparation (PD2 and Spectral Synthesis Task Leader, under guidance of Survey Scientist and ASPCAP Team Leader).
    - Preparation of line list.
    - Calculation of library of synthetic spectra.
- Pre-survey abundance analysis of calibration stars (ASPCAP Software Team members, under guidance of Spectral Synthesis and Abundance Calibration Task Leaders and ASPCAP Team Leader).
    - Before APOGEE spectrograph is on line, reduce currently available, APOGEE-proxy (e.g., ARIES, other) NIR spectra (ASPCAP Software Team members).
    - Acquisition of optical high-resolution spectra, involving proposal writing, data-collection, reduction (ASPCAP Software Team members).

  - Parallel analysis of optical and NIR (ARIES) data, aimed at enforcing consistency between two data sets.
- Stellar parameters pipeline.
  - Outline general methodology for derivation of stellar parameters from combination of photometric colors and APOGEE spectra, using tests on calibration star data as reality checks (Stellar Parameters Task Leader, under guidance of ASPCAP Team Leader and Survey Scientist).
  - Write stellar parameters pipeline code (PD2, under guidance of Stellar Parameters Task Leader, ASPCAP Team Leader and Survey Scientist).
- Chemical abundances pipeline.
  - Outline general methodology for derivation of chemical abundances, using tests on to calibration star data as reality checks (Chemical Abundances Task Leader, under guidance of ASPCAP Team Leader and Survey Scientist).
  - Write chemical abundances pipeline code (PD2, under guidance of Chemical Abundances Task Leader, ASPCAP Team Leader and Survey Scientist).
- Documentation of stellar parameters and chemical abundances pipeline (PD2, proofread and edited by ASPCAP Team Leader, approved by Survey Scientist).

6.3.4. APOGEE Quality Assurance (AQuA) Software

Software will be developed for use by the observers on the mountain for real-time evaluation of APOGEE data quality.

- Decision on whether to write separate code or modify full APOGEE software pipeline software for quick-look option (Survey Scientist, working with all Software Team Leaders).
- Establishment of real-time acceptance criteria for APOGEE observations (Survey Scientist, working with all Team Leaders).
- Writing of AQuA software (PD1, under guidance of ASECS Team Leader; if decision is to adapt full APOGEE software pipeline, then PD2 helps, with guidance of ASPCAP Team Leader).
- Establishment of acceptance criteria for APOGEE reduced data (Survey Scientist working with all Software Team Leaders, the PI and the CPO).

6.3.5. APOGEE Pipeline Coordination

- Manage information flow between different Software Teams and Task Leaders (Survey Scientist).
- Pipeline coordination and interfacing (Survey Scientist).

6.4. Target Selection Implementation

The APOGEE Field and Target Selection (AFTS) Team will define and coordinate the procedures used to select the fields and specific targets within each field that will be observed by APOGEE. This process must take into consideration both scientific and practical considerations. Management of these tasks will be overseen by an AFTS Team Leader, who will act in much the same way as the Software Team Leaders described above.

6.4.1. Definition of Science Targets

The overall scientific strategy for selecting APOGEE targets and field centers will be based on a prioritization scheme after an evaluation of proposed APOGEE science cases. A breakdown of the tasks associated with field and target selection is as follows:

- The PI and Survey Scientist will appoint an AFTS Team Leader, who will be responsible for translating the primary APOGEE survey objectives (as outlined in the original proposal and SRD) into targeting requirements.
- The AFTS Team Leader will also put out a call and be responsible for collecting science cases for potential ancillary science projects that might be pursued simultaneously with, but with minimal impact on, the primary APOGEE survey. The ancillary science projects will be conducted as opportunity permits, but will be of low priority compared with meeting the primary science objectives. Ancillary science cases will spell out the number, type, distribution by position, magnitude, color, and other possible criteria for the targets needed to achieve the goals of that proposed project.
- The AFTS Team Leader will be responsible for organizing and convening a workshop to explain and discuss the proposed plan for executing target selection and field distribution according to the primary APOGEE survey objectives and also to discuss and rank any proposed ancillary science projects. The ranking will be decided by the PI, Survey Scientist and AFTS Team Leader, who will establish criteria for assigning fiber fractions/distributions to any accepted ancillary science programs. Ancillary science that enhances the interpretation of the primary survey science will be given higher weight.
- The AFTS Team Leader will be responsible for elaborating an integrated target selection plan that will outline preferred APOGEE field placements and numbers of visits (i.e. net integration times) for each field, as well as criteria for prioritizing the selecting the number, type, distribution by position, magnitude, color, and other possible criteria, of targets within each field and including targets from any approved ancillary science programs.

6.4.2. Implementation of Target Selection Procedures

Implementation of the science-based targeting strategy must take into account several practical considerations, including coordination of bright time observing with MARVELS and specific restrictions of fiber placement both in the focal plane of the telescope and in the focal plane of the spectrograph.

- The AFTS Team Leader and the Survey Scientist will negotiate and coordinate the APOGEE field and target selection plan with representatives of the MARVELS project.
- The AFTS Team Leader will work with the ASECS Team Leader to formulate a workable and efficient plan for fiber management in the telescope focal plane that will ensure that SRD specifications for cross-contamination of neighboring spectra in the spectrograph focal plane can be met (nominally through a source magnitude binning scheme).
- The AFTS Team Leader will be responsible for merging the science selection, MARVELS coordination and fiber management plans into a working template algorithm for APOGEE targeting. This algorithm will be documented clearly by the AFTS Team Leader.

- Based on the documented AFTS target selection algorithm, code will be developed by PD1 to implement the target selection, under the supervision of the head of the AFTS Team.
- The AFTS Team leader and the Survey Scientist will ensure that the target selection code will achieve a sample of stars that meet the SRD definition of the primary APOGEE science objectives.

6.5. Review & Schedule Milestones

6.5.1. APOGEE Hardware Reviews and Milestones:

- Optics PDR: 2008 Q4
- Opto-mechanics PDR: 2009 Q1
- Cryostat PDR: 2009 Q1
- CDR: 2009 Q2
- Instrument integration and testing: 2010 Q1-Q2
- APO integration and testing: 2010 Q3-Q4
- Science validation: 2011 Q1
- Operation readiness review: 2011 Q2
- APOGEE commissioning: 2011 Q2-Q3

6.5.2. APOGEE Instrument Software Reviews and Milestones:

- PDR: 2009 Q1
- CDR: 2009 Q2
- Acceptance Review: 2011 Q1

6.5.3. APOGEE Software Pipeline Reviews and Milestones:

- Spectral extraction and calibration timeline
  - PDR: 2009 Q2
  - CDR: 2009 Q4
  - Acceptance Review: 2011 Q1
- Radial Velocity and Stellar Parameters and Chemical Abundance Pipelines
  - PDR: 2009 Q3
  - CDR: 2010 Q1
  - Acceptance Review: 2011 Q2

6.5.4. APOGEE Field and Target Selection Review and Milestones:

- AFTS call for science cases: 2009 Q3
- ATFS science case workshop: 2010 Q1
- AFTS plan document: 2010 Q2
- AFTS software written: 2010 Q3

- AFTS Team tests selection software: 2010 Q4

6.6. Commissioning Plan

APOGEE commissioning will entail the following tasks:

- The instrument acceptance tests described in Section 6.7.1 will be undertaken as described there under the responsibility of the personnel named.
- The software pipeline acceptance tests described in Section 6.7.2 will be undertaken as described there under the responsibility of the personnel named.
- The SDSS-III observers will be trained in the operation of the APOGEE instrument, including operation of the instrument software (under direction of the Deputy Instrument Scientist) and the APOGEE Quality Assurance software (under the direction of the Survey Scientist).
- The fiber plate-pluggers will be trained in the procedures for plugging APOGEE fibers, including the APOGEE fiber management plan discussed in Section 6.4.2 (AFTS Team Leader, with the Survey Scientist).
- APOGEE observations will commence in the nominal full survey model for at least one full lunation, including the expected nightly operational throughput of APOGEE field visitations, coordination with MARVELS observing, the Observers successfully using the instrument software and AQuA software to evaluate conditions and APOGEE performance without assistance, and APOGEE data successfully deposited into the archive.
- The Survey Scientist, with the assistance of all Software Team Leaders will process the data from this lunation period and evaluate it against the SRD.
- Once the APOGEE data are verified to be to SRD standards, and the nightly survey operations are deemed to meet the nominal nightly visitation throughput anticipated in the SRD, the Principal Investigator will recommend to the CPO that the APOGEE survey be considered fully commissioned.
- The CPO will approve the assumption of normal APOGEE survey operations.

6.7. Acceptance Tests

6.7.1. Instrument Acceptance Tests

Instrument acceptance tests will be undertaken under the management of the Survey Scientist with test data collected by the instrument team under the management of the Instrument Scientist. The Survey Scientist will be present when these observations are undertaken, to monitor the operation of the instrument and evaluate readiness for routine observations.

The ultimate instrument acceptance tests will be defined in consultation with the CPO. At present, it is envisioned that during APOGEE hardware commissioning, observations will be taken of a series of calibration stars having previously determined, accurately known radial velocities, atmospheric parameters and chemical abundances. Integrations will be taken through different fibers for each star and with integrations scaled to mimic the typical S/N expected for normal APOGEE targets. In addition, test plugplates mimicking typical survey fields will be designed, observed and reduced. The resulting data will be processed by PD1 and PD2 through the AQuA and/or other software that can assess the S/N, throughput as a function of wavelength, spectral range, repeatability/stability, and resolution of the data. Final instrument acceptance will be recommended by the APOGEE PI to the CPO after a report by the Survey Scientist regarding the results of the test observations.

A breakdown of the tasks associated with field and target selection is as follows:

- Calibration test stars for Instrument Acceptance Test selected by AFTS Team Leader. Instrument Acceptance test plate designed by the AFTS Team.
- Test observations described above undertaken by instrument team under management of Instrument Scientist in the presence of the Survey Scientist.
- PI, Survey Scientist and Software Team Leaders, in consultation with the CPO, will decide upon what software/mechanism the instrument acceptance test will be based (e.g., AQuA software reduction of the data, should such software already been accepted, or some other methodology).
- PD1 and PD2 reduce the test data.
- Survey Scientist reports to PI on the results of acceptance testing and makes recommendation for acceptance.
- PI recommends acceptance of the instrument to the CPO.

6.7.2. Software Pipeline Acceptance Tests

Software pipeline acceptance tests will be undertaken under the management of the Survey Scientist with test data collected during science validation. The ultimate software pipeline acceptance tests will be defined in consultation with the CPO. At present, it is envisioned that test exposures of the above-mentioned calibration stars and the survey test plugplate will be passed through all software pipelines, which will need to reproduce the accepted radial velocities, atmospheric parameters and chemical abundances of the standard stars to the accuracy specified in the SRD. The resulting data will be processed by PD1 and PD2. For each software pipeline, the results of the processing will be evaluated by the Team Leader of each Software Team, who will report to the Survey Scientist on the readiness of the specific pipeline under their purview. Acceptance of each piece of the software pipeline will be made by the CPO upon recommendation of the APOGEE PI and Survey Scientist.

A breakdown of the tasks associated with field and target selection is as follows:

- Calibration test stars for Instrument Acceptance Test selected by AFTS Team. Instrument Acceptance test plate designed by the AFTS Team.
- Test observations described above undertaken by instrument team under management of Instrument Scientist in the presence of the Survey Scientist.
- PD1 and PD2 reduce the test data using the standard pipelines (AQuA, ASECS, ARVP, and ASPCAP).
- The Team Leaders of the ASECS, ARVP and ASPCAP Teams evaluate the results of their specific codes and report to Survey Scientist when codes meet SRD specifications. A joint report of all of these Team Leaders will be made for the AQuA software.
- Survey Scientist makes recommendation for acceptance of software pipelines.
- PI recommends acceptance of the software pipeline to the CPO.

7. SEGUE-2 Development

SEGUE-2 is an extension of the SEGUE survey in SDSS-II. It is specifically scoped to use the existing SDSS spectrographs, and will select targets from the final SDSS-II imaging data in DR7. SEGUE-2 has no new software deliverables. Therefore, there is very little development.

7.1. Hardware Development

There is no hardware development in SEGUE-2.

7.2. Instrument Software

To accommodate the increased dynamic range of the SEGUE-2 plates, there is a modification to the existing spectroscopic cartridges in to mark half the fibers in each block of 20 by color. This is used to segregate the brightest fibers on the CCDs so that they are nearest to targets of similar magnitude and well separated from the faint targets and the sky. This is a simple colored sleeve on the fibers.

The allocation of bright, medium and faint fibers was done by Harding, Rockosi and Gunn. The plan and implementation of the marking scheme, which required a special plate, was developed by Gunn and Harding with input from Blanton (plate design), Rockosi and APO plugging technician Diana Holder. Holder did the cartridge mapping. Blanton developed the code that makes the plate overlay files. Testing will be done by Rockosi on the first plate plugged using the new overlays. The test is to check the position of spectroscopic targets within a fiber bundle vs. magnitude.

7.3. Pipeline Software Development

Software development for SEGUE-2 falls into two categories. The first and most straightforward is to the pipeline infrastructure. The pipelines for SEGUE in SDSS-II were run at Fermilab. The SEGUE-2 pipelines will all be run in Princeton. This requires updating the spectroscopic robot so it can get data from the Science Archive Server, run the correct version with the correct inputs, etc. Work to implement QA procedures and to repackage the data as necessary for the data releases also falls in this category. The second category of pipeline development is science driven, with the goal of improving the accuracy, robustness, and/or scientific usability of the SEGUE-2 data.

The infrastructure work will be done by the Princeton software group: Craig Loomis, Steve Bickerton and Fergal Mullally. The work on the SEGUE stellar parameters pipeline will be done by Young Sun Lee at MSU. The work will be done in coordination with the Data Group, led by Blanton. A representative from the Princeton group will attend the weekly data telecons. Their tasks include: implementing and overseeing the SEGUE-2 automated spectroscopic reductions, implementing the data format as agreed upon by the data group and the SEGUE-2 team (represented by Rockosi), working with the data group to get the raw data and reduced flatfiles archived in the SAS and distributed to the collaboration, and getting the data ready to load into the CAS for DR8. In conjunction with Lee, they will get the SSPP under regular configuration management and make any necessary infrastructure changes to get the SSPP incorporated into the automated reductions and archiving.

The science-driven development will be done in coordination with the SEGUE-2 team. Rockosi, Gunn and Harding will be the interface between the software developers and the rest of the team.

The work will be prioritized by the SEGUE-2 team, and will be scoped such that it is finished and tested with time to run the pipelines before the deadline for DR8. The tasks include improving the range of T_{eff} over which the 1-d pipeline returns reliable and accurate radial velocities, incorporating an updated calibration of the outputs of the SSPP and getting good reductions from the SEGUE globular cluster plates.

The infrastructure work will be tested by running the data on all the existing SEGUE and SEGUE-2 plates and checking the outputs against the DR7 (SEGUE) and rerun 102 (SEGUE-2) versions of the reductions.

The science development will be tested by checking the outputs using repeat observations, comparisons with the SEGUE cluster plates, the high resolution follow-up sample, and with the noise model for SEGUE spectra developed by Rockosi. Testing will be planned by the developers with input from Rockosi.

Gunn is the supervisor of the pipeline group in Princeton. The group communicates with Rockosi weekly, either by telecon or with email updates.

7.4. Target Selection Implementation

Target selection for SEGUE-2 was developed by the collaboration scientists. Proposals for how to evolve SEGUE-2 target selection based on what was learned in SEGUE(I) were solicited starting in February, 2008. That discussion resulted in the focus of SEGUE-2 on the distant halo, with more fibers allocated to fainter targets and to target categories that reached into the halo. The changes and additions were decided by consensus within the SEGUE-2 team. The selection algorithms were implemented by Rockosi and checked by the science leads for each category. The infrastructure code that processes the fields was implemented by Rockosi and Harding. All the target selection code has been under configuration management since the beginning, and is documented on the wiki.

The rationale for the different science selection algorithms and the final set of criteria are recorded in the SEGUE-2 target selection document. The information for that document is posted on the wiki and the will be under configuration control for the duration of the survey.

Target selection is run by Gunn and Rockosi, checked by them before being passed to plate design. The plots and other diagnostics are available to be archived, and SEGUE-2 will get that done in coordination with the data group. The science leads for each category checked the initial implementation of the code, and are monitoring the targets as they come in.

7.5. Review & Schedule Milestones

SEGUE-2 passed its readiness review on August 11th, 2008.

The first plates were delivered for drilling July, 2008

The delivery date for the final data reductions and parameter measurements for DR8 is December 1, 2009.

7.6. Commissioning Plan

No hardware or control software commissioning is required.

7.7. Acceptance Tests

The test of the fiber mapping is successful if the fibers are segregated as well as possible given possible mismatch between the average magnitude distribution and that on an individual plate.

The acceptance tests for pipeline infrastructure development are to run the new pipelines on old data and compare all the measurements. The test passes if the values are the same. In situations where differences are expected, the test passes if the expected change occurs.

Acceptance testing of science development work is done in conjunction with the SEGUE-2 team. All the tasks that are put at high priority by the team will have a designated science lead who will work with the developers to design and or do tests, and evaluate the results. Acceptance will be by consensus of the SEGUE-2 team.

8. MARVELS Development

The first MARVELS survey instrument was built at the University of Florida (hereafter UF) in 2007-2008 and commissioned at the SDSS telescope at APO in September 2008. Since September, the instrument has been used for the MARVELS survey. The instrument worked well during the first few months of operation. There was no instrument down time caused by the instrument technical problems. The instrument has produced steady survey data that are being processed at UF. The photon noise limited RV errors are consistent with the forecasted ones. For instance, the photon limited errors for V~8 and V~11 solar type stars in 40 minute exposures are around 5 m/s and ~20 m/s, respectively. The typical peak-to-valley (P-V) temperature variation over ~10 days for the instrument is ~10 mK while the RMS variation is ~3 mK. The remaining major tasks for MARVELS development are the data pipeline and instrument control and interface software development. The current data pipeline still produces substantially higher RV measurement errors than the photon noise limited errors, especially at the faint magnitudes. There are a few survey hardware components that require improving and fixing, including a problem of internal instrument heating caused by the CCD camera shutter and a problem of short fibers breaking.

8.1. Hardware Development

Currently, heat created by the MARVELS CCD camera shutter coil increases the instrument temperature by about 20 mK when the shutter is operated for a long time during the nighttime observations. This temperature increase produced about 100 m/s radial velocity (RV) drifts between the daytime and nighttime when the shutter is closed during the daytime and opened during the nighttime. This RV drift is the single largest drift in the MARVELS instrument over ~10 days. Although this kind of RV drift can be largely calibrated in the RV measurements of science survey targets, this drift has added some uncertainties in the RV measurements on order of tens of m/s. Therefore, the MARVELS top priority for hardware improvement is to fix the shutter heating problem. Several approaches were identified to fix the shutter heating problem: modify the CCD camera control electronics to allow the shutter maintain a constant holding current for both shutter open and close; replace the shutter with a bi-stable shutter which allows the shutter turn off the holding current for both shutter open and close; install a different bi-stable shutter next to the fiber feed location instead.

During the early operation, we found that on average about 1-2 fibers were broken per fiber cartridge after the first two runs (2008 September and October runs). The main cause for fiber breaking is a combination of tight early-drilled fiber holes and aluminum buildup on the ceramic fiber tips. Since then, we worked with the University of Washington machine shop to increase the fiber plughole diameter. By increasing the fiber hole diameter by ~18 μm, fiber plugging becomes much easier. We have been monitoring fiber behavior since to determine if the problem is fixed. At the same time, we are developing a cleaning procedure using 2% NaOH solution. The early testing results show the cleaning can totally remove the aluminum buildup around fiber tips, and we plan to have the APO staff implement a standard cleaning procedure after every 10 fiber pluggings. The team is also developing fiber fixing and maintaining tools for fixing the broken fibers and polishing fiber tips to maintain fiber coupling efficiency after many fiber plugging and unplugging cycles.

Below is the breakdown of the remaining tasks for MARVELS hardware development:

- Shutter heating problem fixing: J. Ge, K. Hanna & F. Varosi
- Installation of fast internet connectors and lightning protection devices: K. Hanna
- Replacement of one motor control end switches: K. Hanna

- Installation of a new ventilation cover: F. Leger & J. Ge
- Fixing of broken fibers: X. Wan
- Fixing short fiber bundles: P. Rohan & L. Chang
- Fiber coupling measurements: P. Rohan & X. Wan
- Fiber fixing and maintenance tools: X. Wan
- Short fiber bundles for fiber mapper: X. Wan & P. Rohan
- Fiber mapper: F. Leger

8.2. Instrument Software

We are currently running the MARVELS instrument using Instrument Control Software (ICS) that is based primarily on a Java program with a few drivers written in C++. Our CCD operation is currently controlled with IDL software. The system is stable and has been running nearly continuously for ~2 months (Nov.-Dec.). Several bugs have been discovered during the first 2 months (Sep.-Oct.) that have since been fixed, but required brief restarts of the software. We have also worked to sync our software with APO software, such as the APO data mirroring software.

Our primary goal remaining is to convert our CCD control software from IDL to C++. Our final software design is to have the software controlled mostly via Java, including a Java front end, with a few C++ drivers, running off a Sun server. Nearly all components are in place, except for switching the CCD control from IDL to C++. The interactive front-end for observers is also in IDL currently, and is a text-based Question & Answer interface. Our main goals, to be implemented in time for the January 2009 observing run, is to finish transitioning the CCD software from IDL to C++ and to finish our Java Graphical User Interface (GUI). The Java GUI will serve as the interface for the observer to take observations and calibrations.

Frank Varosi will be finishing the CCD software conversion from IDL to C++. Scott Fleming will be working on finishing the Java GUI front-end, with assistance from Frank Varosi.

We plan to develop and test the software during December 2008, install in late December and test on-site during the January 2009 observing run, where both Scott Fleming and Frank Varosi will be on-site to monitor observations and fix any problems that should arise. The goal is to have all final software in-place before the February 2009 run begins.

8.3. Pipeline Software Development

The MARVELS preprocessing engine is already in a serviceable form. Based on the data model and FITS header keywords, this engine automatically takes raw science images and calibration frames and creates 1-D extracted spectra ("whirls"), augmented with phase and visibility information. Continued work on the preprocessing will include any adaptations needed to accommodate data format changes as the Instrument Control System, quality assurance, and Telescope Control Computer software suites continue to evolve towards their final form (FITS headers may become richer, and unanticipated format change bugs could creep in as acquisition software changes). The parameter files controlling the preprocessing engine should also be streamlined and converted where possible to “Yanny .par” files or FITS tables. Software version stamps automatically compiled by examining the CVS-controlled source have yet to be implemented.

Anticipated areas for future improvements to the preprocessing engine include further corrections for optical distortion, and improved cosmic ray rejection. Fine tweaking of trim sections and parameters, as well as more intelligent automation, are lower priority tasks, but may improve the results on fine scales. One exploratory aspect that we plan to work on is to rebin the original spectra itself into log-lambda co-ordinates before converting it into whirl form. In this co-ordinate system a Doppler shift is just a linear shift. This may possibly improve the performance of the post-processing software, especially for the low S/N spectra.

The MARVELS postprocessing engine to extract radial velocities from the whirls is currently under redevelopment. This new research and development has become necessary because of the poorer than expected instrument throughput, which causes faint targets to fall into a photon-starved regime the pipeline was not designed to operate under. We anticipate that hardware adjustments currently underway will improve the throughput, but we will continue to modify the pipeline to improve its performance in the lower S/N regime. Also, the plugged fiber to spectrum number mapping as finally installed is more convoluted than planned, which requires new bookkeeping to be introduced. Precision measurements of the interferometer delay from outside the pipeline still need to be incorporated into velocity extraction.

Current avenues of research for improving velocity extraction performance at low signal levels include using high signal template epochs, binning pixels during processing, developing two different extraction codes (referred to as "unwrapping" and "method 14" below), and investigating the amount of ringing during subpixel resampling, as well as attempting to use the 'log-lambda' whirls from preprocessing upgrades.

Next order improvements that will have to be explored after improving faint star velocity extraction include re-calibrating our reduced data with ThAr reference source instead of the iodine absorption to increase our RV measurement wavelengths from ~530Å to ~ 620Å, subtracting contaminating moonlight, and using iodine absorption lines in daytime sky spectra for additional long-term drift monitoring. These improvements require large investments of new effort because they require a degree of spectrum modeling.

Value-added science calculations that should be added to the automated pipeline include stellar parameter calculations including rotational velocities from fringe visibilities, and absolute velocities from comparison of sun-like stellar spectra to day sky spectra.

The current breakdown of outstanding tasks is as follows:

- Preprocessing adaptations to data format changes: B. Lee
- Preprocessing bookkeeping standards cleanup: B. Lee
- Optical distortion corrections: B. Lee
- Preprocessing parameter optimization and added intelligence: B. Lee
- Fiber to spectrum number mapping: N. De Lee and B. Lee
- Postprocessing improvement of signal by better template epochs and/or binning: B. Lee
- "Unwrapping" method experimentation: B. Lee and S. Mahadevan
- "Method 14" experimentation: S. Mahadevan
- Incorporation of interferometer delay: B. Lee
- Ringing in resampling: B. Lee
- Re-calibrating the RV data with ThAr reference measurements: B. Lee & N. De Lee
- Moonlight subtraction: N. De Lee and S. Mahadevan
- Velocity extraction with superposed iodine lines: S. Mahadevan and N. De Lee
- Data archiving and managing software development: N. De Lee and S. Fleming
- Stellar parameters from both SDSS spectra and MARVELS spectra: S. Thirupathi

- Rotational velocities: S. Thirupathi
- Absolute velocities: S. Mahadevan
- Automatic schedule software optimization: B. Lee
- Target and plate tracking software development: N. De Lee & B. Lee
- Data quality insurance software development: S. Fleming, & N. De Lee

8.3.1. Integration and Test Plan

The goal of the pipeline development is to meet or exceed the science requirement RV precisions.

The pipeline will be run at UF on raw data copied from the Science Archive Server. The pipeline will eventually be accessible via the SDSS-III SVN archive, at which point results can be cross-verified by runs at the Science Archive Server.

The data required to test fundamental performance of the pipeline includes a calibration block of biases, flats, arc lamp spectra, and spectra of a tungsten lamp with iodine absorption lines. RV images to be calibrated and extracted should test short timeframes (one image per day for a week) and long timeframes (several images per month for three months). RV images for testing should include at least one faint (V=11-12) and one bright (V<8) reference star of known RV; all RV images should be bracketed before and after by calibration lamp images to track instrument drift.

The rms residual of the RV curve of a reference star compared to its model RV will be compared against the survey requirements for a star of that magnitude to determine whether the pipeline meets goals at low and high signal-to-noise, for short and long time periods. Also, assuming the other stars imaged simultaneously with the reference stars have flat RV curves, they may be used to provide extra data points to check the curve of rms residual as a function of magnitude, with the expectation that the residuals will sometimes be worse because the stars are not known to be RV stable.

The required data are all already in hand. Retroactive runs of the pipeline may be performed as desired. N. De Lee will take a lead in rerunning the data processing.

8.4. Target Selection Implementation

MARVELS target selection begins with photometric selection of FGK stars in the magnitude range V=8-13 carried out by using the GSC2.3 catalog matched to 2MASS. For stars with V<10, a combination of color cuts and reduced proper-motion are used to select stars suitable for MARVELS. For stars fainter than V=10, a color cut of J-K>0.28 is imposed to weed out early type stars unsuitable for precision radial velocity.   Stars that survive this cut are observed with the SDSS spectrographs (or BOSS spectrographs in later years) to reject V>10 giant stars, stars earlier than F8, all of the very active stars based on the CaII emission, and spectroscopic binaries based on the SEDs. The brightest surviving targets, up to the number of fibers in the MARVELS spectrograph(s), are monitored for the main survey. Target selection is detailed in the MARVELS Target Selection document.

The reductions of the spectroscopic pre-selection plates use the IDLSPEC2D pipeline developed at Princeton, with a few modifications to accommodate bright magnitudes, short exposures and low latitude high extinction fields. Stellar parameters are estimated from a weighted mean of the SEGUE Stellar Parameters Pipeline and a method developed at UF, which incorporates more realistic microturbulence values (for the near solar metallicities typical of MARVELS targets), normalized instead of flux calibrated spectra to avoid flux errors in high reddening fields, and alternative estimates of metallicity and log*g*.

The target selection itself uses a custom IDL software program written by Scott Fleming at UF.

The breakdown of target selection tasks is:

- Field selection: Suvrath Mahadevan, Scott Fleming, John Bochanski
- Plates inputs for plate design: Scott Fleming
- Plate design: Mike Blanton, Paul Harding, Demitri Muna
- Observations: APO staff
- Pre-survey reduction and analysis: Sivarani Thirupathi
- Final target selection and plate inputs for RV plates: Scott Fleming.

Several future improvements are anticipated for target selection and/or characterization of stellar targets. Initial starting guesses of stellar types should be based on photometry and isochrones. The synthetic grid with metallicitites above [Fe/H]=0.2, needs to be included, which would improve the metallicity estimations for high metallicity stars. This would not affect the target selection,, since we do not pre-select based on metallicities. The reddening and stellar densities vary among fields, hence in a magnitude limited samples we had a different ratios of dwarfs/giants. In some cases, pre-survey observations of 300 stars (for a double drilled plate) were not enough to select 100 dwarfs/subgiants with log*g* > 3.0. In the future we need to single drill fiber plug plates for high extinction fields.

8.5. Review & Schedule Milestones

The remaining major reviews include the instrument acceptance review and the survey science review. We plan to have the instrument acceptance review in early 2009, likely April or earlier, after we fix the camera shutter-heating problem and also have a fully functional data processing pipeline. Since early survey data (before December 2008) were largely affected by RV drifts caused by shutter heating, we will likely start to accumulate reliable survey data in December 2008 and after. We plan to have the survey science review in July 2009 or earlier. After the survey passes its science requirements, we plan to seek additional funds for a second survey instrument, ET2, which we would hope to have on line in September 2010, and for development of improved final analysis pipelines. Development and review schedules for ET2 will be devised at that point.

8.6. Commissioning Plan

The first MARVELS survey instrument was successfully commissioned at APO in September 2008.

8.7. Acceptance Tests

Prior to shipping to APO, the MARVELS instrument passed tests of individual components and lab tests of the integrated instrument. The remaining tests are to demonstrate its performance on the sky at APO under realistic observing conditions.

The MARVELS Science Requirements Document defines the following radial velocity precision requirements for the ET1 spectrograph:

Table 1. The Baseline Instrument Performance Requirements				
V magnitude	Photon noise limited precision	Moon contamination error	Other systematic errors	Total measurement error
V=8	3.4 m/s	0.3 m/s	13.9 m/s	**14.3 m/s**
V=9	5.4 m/s	0.6 m/s	14.5 m/s	**15.5 m/s**
V=10	8.5 m/s	1.6 m/s	16.1 m/s	**18.3 m/s**
V=11	13.5 m/s	4.0 m/s	19.3 m/s	**23.9 m/s**
V=12	21.3 m/s	10 m/s	25.9 m/s	**35.0 m/s**

Table 8.1: MARVELS ET1 Performance Requirements

The moon contamination error arises because moonlit spectra include a superposed solar spectrum at varying velocities and sky subtraction with an interferometric instrument is non-trivial; this is a requirement on the data pipeline rather than the instrument. We in fact expect to do somewhat better than the requirement on "other systematic errors" because they include a contribution from errors in deconvolving the spectrum of a superposed iodine cell, and we believe that the instrument stability is good enough to allow observations without the cell superposed.

For acceptance of the instrument and the data pipeline, the requirement is that the hardware produce data capable of yielding the rms measurement errors listed in Table 8.1 for 60 objects in a 50-minute exposure and that the data pipeline produce measurements of this precision from the data. The acceptance tests will be based on observations of 4 plates over a 3-month interval, with at least 10 observations per plate:

1. 1 Plate - single bright RV stable reference star. Many observations will allow identification of major error sources and tests of the instrument's fundamental precision limit.
2. 1 Plate - V~8 mag reference star for precision and stability measurements
3. 2 Plates, V=8-12 stars using the full complement of fibers, including at least one known planet-bearing star in each field. These will be observed at several different separations from the moon to allow isolation of moonlight contamination error.

If measurements are not reaching the required precision because of identifiable, and fixable, software limitations, and the data themselves are shown to be good enough to yield the desired precision, then the instrument may be accepted by SDSS-III prior to acceptance of the data pipeline.

To proceed to the final four years of observations and fundraising for ET2, MARVELS must pass an additional science review based on data taken under routine survey conditions. At a high level, this review must demonstrate that the MARVELS target selection, survey strategy, instrument performance, and data pipeline performance, with an additional four years of observations and ET2 performance equivalent to ET1, will achieve the goals set forth in the Science Requirements Document. This review will also consider the scientific viability of continuing MARVELS operations in the absence of ET2.

Criteria to accept the instrument interface and control software

The initially delivered MARVELS interface and control system (ICS) will be a stand-alone system in 2009. It will be integrated into the final telescope instrument control system in late 2009. The criteria for the observatory to accept the stand-alone ICS system are that it:

1. Record all of the required information into the data image head for data processing
2. Operate the instrument effectively without producing instrument and/or data errors
3. Quickly send warning signals to APO staff if there is any hardware operation failure due to the instrument room air conditioner, or calibration motor failures or thermal control.

9. Common Infrastructure Development

9.1. Introduction

Most elements of the mountain infrastructure remain as they were for SDSS-I and SDSS-II, but the new projects and new instruments require some fundamental changes to the system, primarily in the areas of the operations software and, on the hardware front, the whole system of fibers, fiber management, and fiber cartridge handling. Both MARVELS and APOGEE use fixed instruments mounted some distance (a roughly 40 meter fiber run) from the telescope, in a carefully engineered new addition to the plate plugging building built especially for this purpose and finished in the early spring of 2008. Dealing with them requires fundamentally different techniques from the old system with telescope-mounted spectrographs.

The infrastructure work is being coordinated by the Infrastructure Lead, James Gunn, with support by the program manager, Bruce Gillespie, the telescope engineer, French Leger, and a number of engineers distributed across the consortium. The system exclusive of software had a PDR in June of 2008. A critical design review including the software (a separate review) is scheduled for February 2009.

9.2. Production of New Cartridges

The survey in the first two phases used nine cartridges, each carrying 640 3-arcsecond (180 micron) fibers, in two sets of 320 each feeding one of the two spectrographs. The BOSS project will use the existing spectrographs but modified with new CCDs and new gratings for much increased efficiency and somewhat increased wavelength coverage, and will use 1000 2-arcsecond (120 micron) fibers (500 for each spectrograph). These fibers will be terminated at one end with plugging ferrules for insertion into the drilled plugplates and at the other end in blocks that are mounted to the fiber slits that are inserted into the spectrographs when the fiber cartridges are installed on the telescope. Slight modifications to the existing cartridge design will allow us to handle 1000 fibers, but this is at the limit of what is possible. The bright time extension of SEGUE-2 (hereafter SEGUE-2B) will also use these upgraded spectrographs if we are able to raise the funds needed to carry out the survey (see Section 2.3.2).

The `bright-time' programs, APOGEE, MARVELS, and SEGUE-2B, will use, respectively, 300, < 200, and 500 fibers, also for a total of about 1000. All of these fibers are terminated at the plate end with plugging ferrules. The APOGEE and MARVELS fibers will be terminated at the other end not on slit plates, but in multifiber gang connectors which allow the connection of fixed multifiber umbilicals to feed the remotely mounted instruments (in both cases about 40 meters away from the telescope), located in the new addition to the plugging lab. The 500 SEGUE-2B fibers will feed one of the two spectrographs, in an arrangement identical to that for BOSS. It is anticipated that most if not all the time the three bright-time projects will share real estate on a plate and will be working simultaneously.

There was no way for these two fiber systems to coexist in a single cartridge, so it was necessary to fabricate new cartridges (in addition to the nine cartridges that currently exist). The decision was taken that there will be eight BOSS cartridges and eight bright-time cartridges. The cartridges will be identical insofar as possible, with only minor differences to accommodate the multifiber connectors for the bright-time cartridges. The cartridge bodies are cast aluminum, and will be identical (to the modified current cartridges as well) for all SDSS-III cartridges. Eight new castings are being procured from a vendor in Seattle, Washington, and the production is underway as of this writing (December 2008). Seven of these will be machined, with one kept unfinished as a backup spare. There are many small parts in this assembly as well; these will be produced by the physics/astronomy shop at the University of Washington. The production of the new cartridges is under the direction of French Leger and engineer Larry Carey at UW, both of whom also supervised the cartridge production for SDSS-I. The schedule calls for the delivery of the first two fully machined cartridge castings in February of 2009, with one every two weeks thereafter. The smaller machined parts will be delivered on the same schedule. These will be assembled on a schedule that insures that at least five cartridges will be on hand for the beginning of BOSS observing in September 2009; the last BOSS cartridge will be delivered no later than December, though the goal is to have all seven ready by September.

The BOSS cartridges will use the new castings, leaving the current nine cartridges in service for SEGUE-2 and MARVELS observing during the (current) fall 2008-spring 2009 season. When the BOSS cartridges are done, the current cartridges will be taken out of service one at a time for modification to the SDSS-III bright-time cartridges. This will consist of removal of the current 3 arcsecond SEGUE fibers, transfer of the MARVELS fibers (already 2 arcseconds in the current season) to the refurbished cartridges, and installation of the APOGEE and SEGUE-2B fibers. This work will be done at APO under the supervision of French Leger, and will be finished by summer 2010.

The MARVELS fibers are mounted in groups of four in `anchor blocks', which are distributed over the area of the plate. Since there are 200 planned MARVELS fibers, 300 APOGEE fibers, and 500 SEGUE-2B fibers in the new cartridges, a composite anchor block holding 4 MARVELS, 6 APOGEE, and 10 SEGUE-2B fibers will handle 20 fibers, the number chosen for SDSS-I and for BOSS, and makes the mounting and distribution of the 50 anchor blocks identical for the BOSS and bright-time cartridges.

9.3. Procurement and Installation of New Fibers

The first phase of new fiber procurement and installation was done by the MARVELS team. The current MARVELS project uses 120 fibers in two groups of 60. The MARVELS instrument can observe 60 stars at once, and the two groups of fibers allow plugging two fields on one plate, with only the necessity of moving the multifiber gang connector between exposures. The fibers have 1.8 arcsecond (110 micron) cores, with a cladding diameter of 125 microns, so standard communications fiber hardware can be used. The team chose to use 2.5mm ceramic ferrules rather than the 2.1mm stainless steel ferrules used in SDSS-I and for the other SDSS-III experiments, so plate drilling is a bit more complex than for other plates, but this has posed no real difficulty beyond a couple of initial confusion-related mistakes.

Procurement of the BOSS fibers is overseen by Larry Carey and Russell Owen at the University of Washington, and is done in two phases. In the first (completed), prototype 20-fiber bundles were requested from several vendors, with specifications on throughput, finish, dimensions, etc. These bundles when delivered were tested at UW using a fiber tester designed by Owen. The vendor who delivered the best combination of price and performance was chosen (the same vendor, C Technologies, as it happened, who built the SDSS-I fibers). A parallel effort was used to choose a vendor for the custom ferrules for the 190-micron diameter fibers. The first five bundles will be delivered in January 2009 for testing of the production run, and subsequently, beginning in March, at a rate of 50 (one cartridge's worth) per month; this puts the last cartridge's fibers in September, roughly in sync with the delivery of the cartridge bodies. A duplicate of the UW fiber tester is delivered to the vendor so that they can do independent testing (and reporting) of the bundles as they are completed. Spot tests will be done after delivery to verify the reports from the vendor.

The SEGUE-2B fibers, identical in every respect to the BOSS fibers, and the APOGEE fibers, which are mechanically identical to the BOSS fibers but are made of ‘dry’ IR-transmissive glass and are terminated in connectors instead of the slit-plate blocks, will be ordered as an extension of the BOSS contract, and will be in hand for the conversion of the present cartridges to SDSS-III bright-time cartridges beginning at the start of CY 2010. Again, this work will be done at APO under the direction of French Leger.

The multifiber gang connectors for the umbilicals were a major development effort also under the direction of Leger, but is finished and works well. A connector can handle up to 10 USConec multifiber connector blocks, which allows plugging and unplugging all of the 300 APOGEE fibers or either set of 60 or 120 MARVELS fibers with one simple operation.

Maintenance of the cartridges is routine. Fiber bundles are replaced when two fibers in a bundle of 20 are broken. This has happened only twice in the 8-year history of SDSS.

9.4. Cartridge Handling Mechanisms

The fiber cartridges are currently handled by specialized and complex mechanisms designed and built by Fermilab for SDSS-I. These machines allow up to nine cartridges to be manipulated within the fiber plugging laboratory, placed in a special handling lift/rack with nine slots which has an opening to the outside of the building, and an identical machine on the exterior to remove the cartridges one by one from the rack and place them on a special cart for transport to the telescope. Within the lab, the manipulator can place and remove a cartridge to/from another handler, the plugging station, which can safely turn the cartridge over for plugging and allow two plugging technicians to work on the cartridge simultaneously. The system has worked well, but it cannot handle more than nine cartridges, and it poses a serious maintenance problem.

In the spring of 2009 a new system will be built for installation that summer. The interior handler and plate plugging station will be removed, and will be replaced by a modified industrial walker/stacker. This is a small electric vehicle which is intended for use in warehouses and retail outlets to handle boxed merchandise--in effect, a miniature fork-lift. The normal lifting fixture will be removed and replaced by a special table to which a cartridge can be clamped. The cartridge can be rotated about an axis parallel to the plate through its center of gravity to invert it for plugging, and 90 degrees about the optical axis for access to the plugging openings. The machine will thus replace both the manipulator and the plugging station, and it is freely mobile, so it can be used to move cartridges from fixed racks in the laboratory to/from the nine-cartridge lift rack with access to outside.

A set of mechanical/electrical interlocks will be designed into the system to prevent damage to the cartridge when loading or unloading racks and to prevent any but very low-speed transport when the cartridge is at any but its lowest height to alleviate any tipping danger.

Another specialized fixture for the machine will be designed and built to facilitate the installation of fibers into a cartridge and maintenance. This fixture will be mounted to a spare walker/stacker at the University of Washington during the cartridge construction.

The design work for these machines is being done by French Leger and Nick McDonald at APO, and Leger will supervise the construction, which is anticipated to be done mostly by CVE Machine in High Rolls, only a few miles from the observatory.

The designs for these machines existed only in preliminary form at the Infrastructure PDR in June of 2008; the CDR has been delayed till February mostly to allow time for this work to be finalized, since there are important safety issues in the design and implementation of the interlocks.

Maintenance of the new machines will be undertaken as a routine task for the APO crew. We have, of course, no history on the machines, but the use is light compared to normal industrial/commercial applications, the loads are small compared to capacity, and we expect no major issues.

9.5. Fiber Management and Plate Marking

Another major change in the handling of fibers is the issue of fiber management in plugging. In SDSS-I,-II, there was no plugging management: any fiber could be plugged into any hole, and the identity of the fiber in a given hole is determined by an automated plate mapping machine. In SDSS-III this is generally not the case; the different projects require differing degrees of fiber management, ranging from specifying which hole a given fiber goes in for MARVELS to no management at all for APOGEE. In addition, it is planned that the bright-time plates will be plugged for the three bright-time projects and they will share the focal plane, so it is necessary to have some way to distinguish the fibers for different projects. This will be accomplished by a color-coding scheme that has been implemented and tested. The fibers are individually coded in groups which correspond to the fibers in a given anchor block for each project, and the plates will either be marked by hand using `overlays' and a digital projector or will be marked by a commercial firm in Seattle using a large flatbed plotter designed for painting signs. Both have been demonstrated to work; the choice will be made on economic and efficiency grounds, this spring, after we have had a chance to evaluate both options somewhat more thoroughly.

9.6. Guide Camera

The current guide camera, an inexpensive commercial device using a very small, uncooled CCD has never been completely satisfactory. Especially going to the new smaller fibers, it was felt that it should be replaced with one with greater efficiency and lower noise, so that, in particular, the sky is well registered on the guide fibers around the guide stars--a major problem with the old one was that flexure in the rather flimsy camera and in the coupling lenses results directly in guiding errors, because the fibers themselves are not seen in the guide exposures--only the stars are seen and are servoed to the centers of fiber images made by use of a lamp at the beginning of the exposure.

A new camera using a cooled, back-illuminated E2V 1Kx1K 13.5-micron pixel device has been procured from APOGEE Instruments (no relation to the APOGEE survey). The camera has an ethernet interface with control software identical to a similar camera that has been in use on the APO 3.5m telescope for some years. Paul Harding at Case has been the responsible person for this project; the camera is finished and has been tested on the telescope, and some preliminary software to interface with the present control system has been written. It will be tested this winter when the press of other tasks (Harding is now handing off the full-time job of plate design to Demitri Muna at NYU) has lessened.

A completely new top end for this software will be written for the new operations software system, discussed below.

The similar guide camera on the 3.5-meter has been trouble-free, and we do not expect trouble from this one.

9.7. Telescope Operations Software

The operations software for SDSS-I and SDSS-II was written specifically for the spectrograph/imaging camera system in use for the earlier phases. It was written mostly in Tcl/TK, is not very modular, would be very difficult to adapt to multiple instruments in use simultaneously, and already before SDSS-III has proven very difficult to maintain and modify. It was deemed extremely inadvisable to attempt to extend it to cope with the much larger scope and new instruments for SDSS-III.

A much more attractive option is to build on the new, modular control system recently written for the APO 3.5m telescope (TUI/TRON) in Python by Russell Owen and Craig Loomis. This task is underway now. It is being jointly managed by Gunn as Infrastructure Lead and David Kirkby from UC Irvine, capitalizing on the latter's experience with the BABAR control system. The programming team consists of Kirkby, Owen, Loomis, and Robert Lupton, who was responsible for a substantial part of the infrastructure for the current SDSS system, most of which will remain in place for the new system.

The system consists of several modules called Actors, Commanders, the Hub, and the Archiver. The commanders (TUI in the previous implementation) are responsible for generating commands in a common format and for the general user interfaces. All commands and replies thereto flow through the hub, which acts as a communication center. The Actors are control systems for instruments, the telescope, the guider, the common scientific data archive, and passive systems like the weather station and the DIMM. The archiver is a special Actor, a front end for a MySQL database that keeps track of all commands, replies, and streams of engineering data from the telescope. Kirkby is writing the archiver, Loomis the hub (he wrote the one for TUI/TRON), and Owen, who wrote most of the TUI code, most of the actors, though three will be written by the instrument teams (BOSS, MARVELS, and APOGEE) in consultation with the ops software team.

At this point, all of the protocols have been defined, and the writing is underway. The schedule calls for delivery of the hub in January of 2009, and the staged delivery of the archiver, which is the most complex single component, finishing in April of 2009, with the system ready for testing as a a complete system in the middle of May. Operations using it will begin in September, as BOSS comes on line.

A critical design review of the system is planned for February 2009.

Maintenance of the system will be handled by Kirkby and others at UC Irvine, with input as necessary from the other programmers.

10. Mountaintop Operations and Maintenance

Mountaintop Operations includes the Observing Systems and Observatory Operations at Apache Point Observatory (APO), Sunspot New Mexico. Key aspects of day and night operations and maintenance, and data acquisition are described in the following sections.

10.1. Observing Systems

The Observing Systems Lead (Jim Gunn) directs the work of the Observing Systems organization and is responsible for ensuring that the observing equipment and systems meet the science requirements and operational needs of the survey. A preventive maintenance program has been established to ensure system reliability. A spares inventory will be maintained through the end of survey operations to ensure system availability. A configuration management program has been implemented to manage improvements in a controlled manner and ensure maintainability. A quality assurance program tracks system performance. The SDSS Problem-Reporting Database is used to report and track concerns and equipment and software problems.

10.1.1. Instrumentation

During development for each instrument the Telescope Engineer and Lead Observer participate in coordinating instrument integration at the observatory. The Telescope Engineer in conjunction with the instrument team verifies that installation, power, and data requirements, and operational and support needs are coordinated with the observatory prior to instrument delivery. In addition, handling of the instrument upon delivery and space requirements for testing and integration are finalized as it nears delivery. The Lead Observer participates in instrument installation, software development and operations integration, and documentation and training for the observers. Upon completion of commissioning, which is accomplished by the instrument team and the staff observers, documentation and spares as well as any specialized training and tooling are incorporated into observatory operations. This incorporation is required to formally accept the instrument.

Following acceptance, the observatory staff assumes responsibility for operation and normal maintenance of the instrument. Minor upgrades may also be accomplished by the observatory staff with coordination and review by the instrument development team. This work may include integration upgrades, out of dewar electronic and mechanical repair, and vacuum and cooling maintenance. Inside dewar repairs to optics, electronics, vacuum, and cooling are generally left to the instrument builders to support until the instrument is no longer needed for the survey. Instrument Control Computer (ICC) software is maintained by the instrument development team.

10.1.2. Plug Plate Operations

The fiber plugging cycle begins after the drilled plates are received at APO. The plates are inspected for shipping damage and cleanliness by the APO plugging staff, logged in the plug plate database, marked for particular fibers or fiber bundles, and stored in the staging facility until they are ready to be plugged. Each morning during an observing run, the plugging staff receives instructions from the observers denoting which plates were successfully observed during the previous night and can be unplugged, and designating which plates need to be plugged for the coming night. When the observations for a plate have been declared complete, the observed plate is unplugged, removed from the fiber cartridge, and returned to the appropriate storage system. New plates are retrieved from the storage system, loaded into fiber cartridges, and plugged with optical fibers. Once plugged, the BOSS and SEGUE cartridge/plate assembly is scanned by the Fiber Mapper, which correlates fiber number with position on the plate. MARVELS fibers are plugged per a pre-determined overlay.

After the first year of observing, a verification process for MARVELS can be run using a slit head designed for mapping of the APOGEE fibers. Once mapping is complete, the cartridge is placed in the storage rack in preparation for spectroscopic observation. Relevant information is loaded in the SDSS plug-plate database throughout the fabrication, plugging, and storage process. APO maintains on-site storage and staging facilities for approximately 300 plates and contracts for off-site facilities for long-term storage of the remaining plates. All plates are stored until the completion of the survey.

Initially nine cartridges will be in use, growing to 16 when BOSS begins survey operations. Existing cartridges will require modification to support fibers for the new instruments. Modifications are done at the observatory when the cartridge is not needed for observing. Due to the increase in cartridges a new storage area and cartridge handling/plugging station is required and will be installed in the existing support building. The installation and integration of these systems will be done by the observatory staff upon delivery.

The plug plate computer hardware and database software is maintained by the observatory staff.

10.1.3. Observers' Programs

Observing Systems includes the Observers' Programs, which are used by the observers to control the telescopes and instruments, as well as the onsite quality assurance programs and spectroscopy planning database. The Observers' Programs include the Imager Observer's Program (IOP), the Spectroscopic Observers' Program (SOP) and the MARVELS instrument control software. IOP provides the observer's interface to the systems used in imaging operations (telescope, imaging camera, and DAQ). SOP provides the observer's interface to the systems used in spectroscopic operations (telescope, multi-fiber spectrographs, and DAQ). The MARVELS software controls their spectrograph and DAQ. Quality assurance software is also available to verify that the data taken meets survey standards. Site software also includes the plate database which is used for planning spectroscopic operattons. The Observers' Programs are developed by SDSS programmers from different institutions and are commissioned and maintained by the Observatory staff.

10.1.4. Data Acquisition System

Observing Systems includes the Data Acquisition System (DAQ) at APO, which consists of the hardware and software that collects data from the instruments and is subsequently sent offsite via a file server and high speed internet link for data processing. Initially the DAQ system, including spares, consists of the SDSS II DAQ hardware and software. Spares for this system are available except for a few fiber optic interface boards that would have to be repaired if they fail.

Early in the program a new DA will be developed. Observatory Information Technology (IT) personnel are involved in the development of this new DAQ. Upon integration at the site and successful commissioning the observatory staff will assume responsibility for hardware maintenance. A spares plan will be established for the new DAQ system and spare components procured to ensure that an adequate supply of critical parts is available to keep data acquisition activities on-line. APO is responsible for specifying and maintaining the on-site spares inventory. Software maintenance of the new DAQ will be split between observatory staff (Stauffer and Brinkmann), University of Washington (Owen) for the Telescope Users Interface (TUI), Princeton (Loomis) for the hub software, and UC Irvine (Kirkby) for the archiver and overall maintenance management.

10.2. Observatory Operations

Observatory Operations captures the work and costs associated with staffing and maintaining the observatory at the level required to sustain survey operations over the 6-year survey period. Observing is accomplished every night of the year unless weather or telescope repairs prohibit it. The organization chart for Observatory Operations is shown in Figure 10.1.

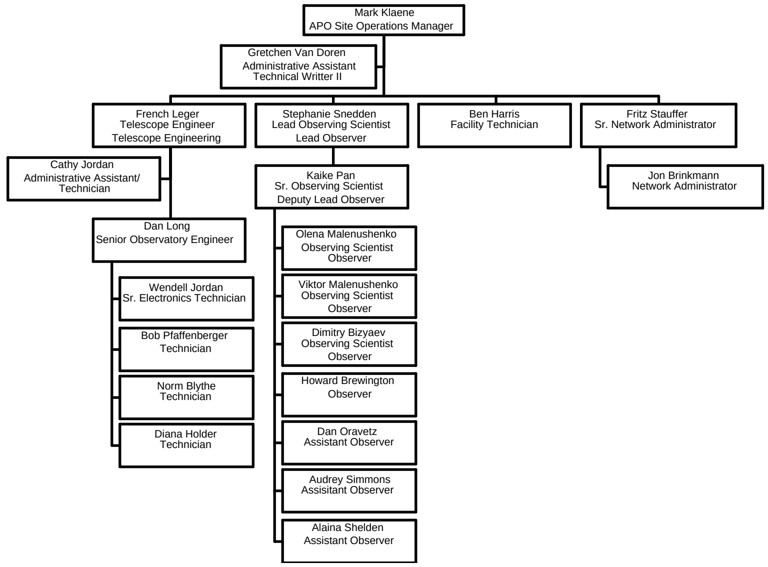

Figure 10.1. Organization Chart for Observatory Operations

10.2.1. Management

APO site management provides managerial and administrative support for observatory operations. The APO Site Operations Manager is responsible for day to day operations as well as staff recruitment and training, budget, and procurements. These activities are developed in concurrence with the SDSS Director and Program Manager to ensure that SDSS funding and scientific goals are met. New Mexico State University (NMSU) is the fiscal agent for ARC regarding mountaintop operations. Administrative personnel both at the observatory and at NMSU perform purchasing, resource management, accounting and inventory control in accordance to ARC and NMSU policies as stipulated under the terms of the grant. The Site Operations Manager provides routine verbal and written quarterly status reports to the Project Manager and has implemented programs for publishing night and day site activity logs which are widely distributed throughout the project.

10.2.2. Observing

APO provides a staff of Observers for the 2.5-m telescope. One observer has been appointed the Lead Observer (Stephanie Snedden) and another Deputy Lead Observer (Kaike Pan). The Lead Observer is responsible for preparing the schedule and activities of the observing staff including maintaining the observers' programs. The Lead Observer is also responsible for interfacing with the engineering and observatory support staff to ensure good communication between management, technical, and observing personnel. The Lead Observer works with the Survey Coordinator to prepare monthly observing plans and is responsible for the execution of those plans. The Deputy Lead Observer assists the Lead Observer in these responsibilities and assumes them in the Lead Observer's absence.

The observing staff is responsible for the safe and efficient use of the telescopes and instruments to collect data that meet the survey requirements. The Observers are tasked with determining which objects/plates and program to observe at any given time based on scheduling, availability, weather, telescope and instrument availability, and priority. These conditions change on the order of nightly to hourly. At the end of each night they must determine the successfulness of the observing so that plans can be made by the daytime crew to prepare for the next evening. During observing they also are tasked with the safety of the telescope, optics, instruments, and personnel. They have the ultimate responsibility to ensure that operations are done to protect personnel and equipment while providing optimum data acquisition. To perform this, they develop and use documented procedures that promote safe and efficient operations, troubleshooting, data quality and uniformity; and they implement and use performance metrics to track observing efficiency and performance against pre-set goals. Short, on-line QA analyses are done to verify data quality; data passing these tests are forwarded for processing. The observing staff publishes observing logs to document observing activities and records problems in the SDSS Problem-Reporting Database.

The Observers also maintain the instruments that monitor seeing and night sky conditions, including the Cloud Camera, and the DIMM. Other Observers maintain software and databases used in support of observing.

10.2.3. Engineering

APO provides technical personnel for telescope engineering in the preparation and maintenance of the 2.5-meter for daily operation. Much of this work is under the guidance of the telescope engineer. The Senior Operations Engineer directs and oversees the day-to-day tasks by this staff. The staff performs cleaning of the optics under detailed procedures and on regular timelines. They perform routine preventive maintenance every month or as specified by a predetermined schedule. Repairs are conducted immediately if required for data acquisition or at the next available time based on science priorities. The plug plate team is included in these staff. APO staff also maintain the meteorological instruments and the engineering and guide camera.

Annually the telescope and instruments are shut down during the summer rainy season (typically for about one month) for general maintenance and overhaul. Activities during this time include major repairs or improvements that take more than 1-2 days to perform, mirror re-coatings and cleanings, and servicing of systems that require optics removal. This activity lasts about 4 to 6 weeks dependent upon tasks, and it often utilizes the additional support from other partner members. These outside engineers and scientist are also occasionally called upon if the issues exceed the expertise of the staff. This is most common with problems inside the dewars of the instruments.

When deficiencies are identified on the telescope, plate plugging, or observatory support functions, the staff will investigate and propose solutions for improvements and upgrades. The Management Committee reviews these plans and, depending on the complexity of the systems, detailed reviews may be held prior to machining or purchasing parts. Installation is coordinated with project management to avoid disruption of data collection.

10.2.4. Support

APO maintains the ARC-provided real property and equipment necessary to support data collection activities at the observatory including office space, laboratories, plug plate lab, physical plant, LN2 servicing equipment, power, and storage. APO maintains ARC-provided facilities and equipment necessary to operate and maintain the instruments, such as vacuum equipment and a Class-100 clean room for maintenance and repair activities associated with the scientific instruments. APO maintains the site telecommunication system and all on-site SDSS computer systems and associated spare parts. These consist of the DAQ systems as well as ARC-provided instrument and telescope control computers. APO also maintains the basic site services and facilities necessary to carry out operations in an efficient manner. Critical equipment and sub-systems have been placed on Uninterruptible Power Supplies, and multiple emergency power generators are available.

The APO operations staff maintains the site facilities, roads, grounds, and housing used by staff and visitors working on SDSS activities at APO. A small machine shop at APO is outfitted with tools and equipment suitable for performing small machining jobs. Larger machining jobs are done under contract with local shops.

11. Survey Coordination

Survey Coordination captures the work associated with planning and executing the observing strategy for the Survey. This includes the development of strategy planning tools that are used to determine the optimum time(s) when specific regions of the survey area should be imaged and when specific spectroscopic plug plates should be exposed. These tools are critical to achieve efficient observing operations and to meet survey performance goals.

The Survey Coordinator (Don Schneider) is responsible for creating detailed monthly observing plans for the Observers based on and prioritized according to the plan described in Section 2.5. The Coordinator is responsible for tracking survey progress against the detailed metrics for each of the four surveys (Section 2.7 and Appendix A) and adjusting observing priorities as necessary to ensure that the survey goals are achieved. This plan also incorporates the planned maintenance (including an expected month-long summer shutdown each year); this will allow efficient planning and scheduling of resources and the identification of periods during which engineering and maintenance work can be integrated with observing activities.

The Survey Coordinator is responsible for developing and implementing survey planning and strategy tools to ensure efficient survey operations and measure observing performance. The Coordinator is responsible for coordinating the steps in the production of plug plates, from target selection to plate inventory at APO. This includes scheduling and oversight of target selection, determining the number of plates required per drilling run, overseeing the generation and delivery of drill files to the University of Washington machine shop, and overseeing plate delivery to APO. The Survey Coordinator is responsible for maintaining the on-line plug plate database, which tracks the location and observing status of all plug plates.

The Survey Coordinator chairs a weekly telecon that reviews the Operational Status of the survey. Attendees of this telecon include the representatives from each of the four surveys, observers, data archive and drilling teams, and APO engineering and management. The Survey Coordinator is responsible for seeing that operational problems are addressed in a timely manner.

The Lead Observer serves as the Deputy Survey Coordinator and is responsible for implementing the monthly observing plans at APO. The Deputy Coordinator is responsible for developing the planning and tracking tools necessary to efficiently conduct observing operations at APO.

12. Plate Drilling

Overall responsibility for the Plate Drilling process lies with the Survey Coordinator. As stated in the Survey Coordination section, the Coordinator is responsible for coordinating the steps in the production of plug plates, from target selection to plate inventory at APO. This includes scheduling and oversight of target selection, determining the number of plates required per drilling run, overseeing the generation and delivery of drill files to the University of Washington machine shop, and overseeing plate delivery to APO. The Survey Coordinator is responsible for maintaining the on-line plug plate database, which tracks the location and observing status of all plug plates. It is expected that there will be a plate drilling order approximately once/month.

The logistics of the plate drilling are executed by the Drilling Coordinator. Each of the survey teams (BOSS, SEGUE, MARVELS, APOGEE) designates one of their members to be the drilling liaison with the Drilling Coordinator. The team liaisons submit their plate design requests to the Drilling Coordinator according to the schedule developed by the Survey Coordinator. The requests include proposed target lists, priorities, hour angles, and whether specific fiber/hole pairing must be respected for each observation.

The Drilling Coordinator, in consultation with the Survey Coordinator, reviews this information and creates the computer numerical control instructions for the drilling machines that produce the plates, files for measurement and quality control tests, plug-mapping files to be read by the plate mapper, overlay files to be used as guidance for the pluggers, and appropriate meta-data for use by the APO plate database. The Survey Coordinator, in consultation with the SDSS-III University of Washington shop lead, monitors the production of the plates and their delivery to APO.

13. Data Processing, Archiving, and Distribution

The previous development chapters have described the plans for pipeline development. This chapter describes the procedures by which the pipelines get run and data get archived and distributed, especially who is responsible for what.

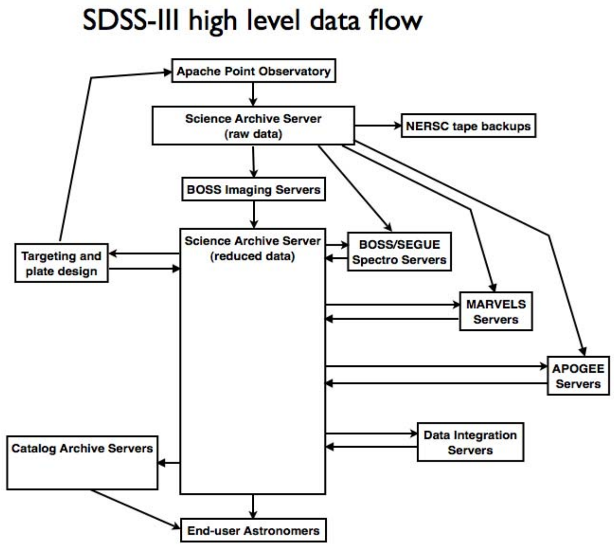

Figure 13.1: SDSS-III high level data flow diagram.

13.1. Data Archiving

The high-level data flow is shown in Figure 13.1 above. During observations, imaging, spectroscopic, and meta-data are written in real time stored temporarily in a staging area on an MJD basis. Checksums are produced at the time of observation. They are transferred to the SAS at LBL via high-speed internet connection. MJDs will only be deleted sometime after they have been archived and backed up to tape; a page on http://trac.sdss3.org conveys the status of each MJD for this purpose. Simultaneously, a second copy of the raw data is stored onto removable disks that are ultimately kept in offline storage at APO. The APO systems administration team (Jon Brinkmann and Fritz Stauffer) ensures that these steps are taken, in consultation with the instrument scientists.

Each morning after data for an MJD is taken, SAS retrieves the raw data and verifies its contents, checking for consistency among the files. It then backs the MJD up to the NERSC HPSS tape system. After each step (copy, verify, backup) SAS updates an internal database with the MJD status. A page on http://trac.sdss3.org posts this status, indicating after the backup step that the MJD is "ready to delete." In the evenings (APO time) the Science Archive Mirror (SAM) at NYU will update its mirror of SAS. The integrity of the data is checked through periodic checksums. The Data Archive Scientist is responsible for this process.

The individual teams download raw data from the SAS. In cases where the reduction pipelines depend on external datasets, the version of those external datasets used (2MASS, Tycho, GSC 2.3, SFD, etc) will be stored at the SAS and mirrored to the reduction institution, so that the pipeline codes have a reliable, consistent, and documented source for their inputs. In consultation with liaisons from each team, the Data Archive Scientist is responsible for maintaining these data sets and coordinating the downloads.

After data processing is complete for any portion of data (an imaging run or a particular instrument on a spectroscopic plate), the SAS will retrieve the reduced data. To facilitate this process, each reduction team will maintain a "rerun file" indicating the status of each re-processing of the data. Only those re-processings marked "complete" will be retrieved. Each set of reductions will be marked with version numbers indicating the software used. The survey teams will maintain a copy of their reductions, though not necessarily a copy of the raw data. When a re-processing is retrieved, it will be backed up in the HPSS system and also copied to the SAM. The upload and backup procedure will be maintained by the Data Archive Scientist.

Because of the timing of SEGUE-2 in the first year, it does not have access to the final reprocessing of the SDSS imaging data. Thus, for targeting and spectrophotometric calibration it relies on SDSS-II DR7 imaging catalogs. The relevant catalog entries used are stored in standard formats tracked in the data model with backups archived at SAS, SAM and HPSS.

Results from data assembly steps occurring at the SAS (window function determination, matching of spectra to photometry) will also have tracked versions, and be backed up at SAM and HPSS. The Data Archive Scientist is responsible for these backups.

Allowed file types as outputs are FITS, FTCL parameter files, CSV files for database loading, and XML files. For each set of data that is produced (raw data, meta-data, reduced data, integration data) there will be a "data model": HTML files that describe for each type of file the directory location, naming convention, required and optional header keywords, required and optional column names, and any relevant data types, units, and descriptions for each. The data model documentation is stored in an SVN product ("sas") and a periodically updated version distributed on the web site http://sdss3.org/internal/datamodel. Changes to the data model must be approved by the Data Coordinator.

The total archival needs for SDSS-III are estimated to be 100 Tbytes. About 50% is raw data and 50% reduced. About 70 Tbytes are required by July 2009, and the needs grow by 6 Tbytes/year thereafter. We assume here that APOGEE does not require Fowler sampling, which would increase the raw data storage needs. The full breakdown is shown in Table 13.1 below. Final calibrated catalog information (about 15 Tbytes) should be on robust servers with 100% up-time; raw data and spectra will be on commodity hardware (with multiple backups to handle failures smoothly).

SDSS-III Data Volume	July 1 2008	July 1 2009	July 1 2010	July 1 2011	July 1 2012	July 1 2013	July 1 2014	
raw								
SDSS-I, -II spec	2	0	0	0	0	0	0	**2**
SEGUE-2	1	0	0	0	0	0	0	**1**
SEGUE-bright	0	0.5	0.5	0.5	0.5	0.5	0	**2.5**
BOSS spec	0	1	1	1	1	1	0	**5**
MARVELS	1	1	1	1	1	1	0	**6**
SDSS-I, -II imaging	28	0	0	0	0	0	0	**28**
BOSS imaging	2	0	0	0	0	0	0	**2**
APOGEE	0	0	1	1	1	1	1	**5**
reduced								
SDSS-I, -II spec	1	0	0	0	0	0	0	**1**
SEGUE-2	1	0	0	0	0	0	0	**1**
SEGUE-bright	0	0.5	0.5	0.5	0.5	0.5	0.5	**3**
BOSS spec	0	0.5	0.5	0.5	0.5	0.5	0.5	**3**
MARVELS	0.5	0.5	1	1	1	1	1	**6**
SDSS-I, -II image cats	14	0	0	0	0	0	0	**14**
BOSS imaging cats	2	0	0	0	0	0	0	**2**
BOSS datasweeps	1	0	0	0	0	0	0	**1**
Imaging frames	0	10	0	0	0	0	0	**10**
APOGEE	0	0	0.5	0.5	0.5	0.5	0.5	**2.5**
Data integration	2	1	0.5	0.5	0.5	0.5	0.5	**5.5**
increment	55.5	15	6.5	6.5	6.5	6.5	4	
total	55.5	70.5	77	83.5	90	96.5	100.5	**100.5**
increment (stable)	5	2	1.5	1.5	1.5	1.5	1.5	**14.5**
total (stable)	5	7	8.5	10	11.5	13	14.5	14.5

Table **13.1**: SDSS-III data volume needs, broken down by survey and data type. Each year's increment is listed. Some data we need to keep on very robust servers (rather than commodity hardware). Out of the full 100 Tbyte, these data account for about 15 Tbyte, which we track in the final rows.

13.2. Data Processing

Data processing activities include all work associated with the development and maintenance of the data processing pipelines, and the organization, operation, and maintenance of the data processing factories.

As described in Section 1.8, SDSS-III data processing is divided among several institutions: LBL for BOSS imaging and spectroscopy; Princeton for SEGUE-2 spectroscopy; University of Florida for MARVELS spectroscopy; and University of Virginia for APOGEE spectroscopy. The PI for each team is responsible for each data processing unit. On-mountain processing occurs at APO for quality assurance purposes. The Lead Observer is responsible for the on-mountain processing. Data assembly steps such as window function determination and matching of imaging to spectra occur at SAS. The Data Coordinator is responsible for the data assembly and for setting the overall schedule for reduction completion.

BOSS imaging reductions will produce a set of photometric catalogs, astrometric and photometric calibrations, datasweep files, and targeting and tiling results.

BOSS spectroscopic reductions will produce calibrated spectra for each fiber, calibration results for each fiber, spectroscopic parameters associated with these spectra and quality flags on each spectrum.

SEGUE-2 spectroscopic reductions will produced calibrated spectra for each fiber, calibration results for each fiber and spectroscopic parameters associated with each spectrum. The SEGUE-2 team will also produce the targeting information associated with each field.

MARVELS spectroscopic reductions will produce calibrated spectra and spectroscopic parameters for pre-selection plates observed with the SDSS or BOSS spectrographs, and extracted results for each fiber along with radial velocities from the MARVELS spectrograph. The MARVELS team will also produce the targeting information associated with each field.

APOGEE spectroscopic reductions will produce wavelength-calibrated, telluric-absorption corrected spectra for each observation, co-added spectra for multiple observations of the same star,, and associated spectroscopic parameters (radial velocity, atmospheric parameters, abundances, and uncertainties). The APOGEE team will also produce the targeting information associated with each field.

For each pipeline that is run, the outputs will track the version of the pipeline and the reprocessing number (as we describe in Section 15.3, the versioning and dependency trees are tracked with SVN and ExtUPS).

The data assembly step is special in the sense that it incorporates multiple survey inputs and produces science-ready files, with associated statistical weight maps and random points where appropriate. The Data Coordinator is responsible for this processing step, relying on help as appropriate from the science teams. This assembly also includes matching objects across the various surveys.

Data assembly depends on input from external data sets, in particular the SDSS-I and SDSS-II surveys and geometry files. The Data Archive Scientist will be responsible for maintaining all external data sets except for the SDSS-I and -II legacy data, which are the responsibility of the Data Coordinator.

Where surveys depend on reductions of other surveys' data (such as APOGEE on SEGUE-2, or BOSS spectroscopy on BOSS imaging, or MARVELS on pre-selection plates from SDSS) they will copy the necessary reductions from SAS.

Data reduction quality is checked through a series of checks described in Chapter 14.

13.3. Data Distribution

The data distribution for SDSS-III is performed through the Science Archive Server at LBL and the Catalog Archive Server at JHU. Mirrors of the CAS will be run at several international locations (Europe, Asia, and South America) to provide fast local access world-wide. A Science Archive Mirror exists at NYU to mirror the LBL site, but it is not expected to be widely used for distribution.

The SAS is run at LBL and maintained by the Scientific Cluster Support team there. It will allow access through an HTTP interface (e.g. available to web browsers, wget and curl) as well as read-only rsync support. A web front-end will allow simple search and retrieval operations, in particular for images and spectra. Private data will be available to collaborators through a password-protected interface. Public data will be available without password protection. The Data Manager will oversee the development of this web site.

The SAS will be mirrored offsite at the SAM (at NYU). The Data Archive Scientist will maintain this mirror.

The CAS is run at JHU and is maintained by the Database Development Group there. It allows access through a MicroSoft SQLServer interface that allows for very flexible querying. An associated CASJobs interface allows server-side caching of and operations on intermediate results. A lay-level front-end called SkyServer distributes navigable JPEG images of the sky, with labeling referring back to the full database and GIF images of spectra for visual inspection. Private data will be available to collaborators through a password-protected interface. Public data will be available without password protection. The CAS Head oversees the development of this database. One dedicated FTE (possibly split across two individuals) will administer the database.

The CAS will be mirrored at multiple other sites, to be determined. The main database, but not CASJobs or SkyServer, will be required to be included in the mirror. Contacts at those sites will administer those mirrors, and coordinate update with the CAS Head. Mirror sites may run CASJobs are SkyServer if desired, but ARC does not take responsibility for maintaining those sites.

Loading of the database will occur in a staged manner. Calibrated FITS files will be created on SAS that are as nearly parallel as possible with the final CAS tables. These FITS files will be converted to temporary CSV files that will be transferred to JHU (either on portable disk storage or over the network) for loading into the database. CAS requires some internal calculations to occur before it is ready for consumption.

As close to 100% uptime as possible will be maintained on SDSS-III CAS and SAS servers serving the current data release to the public.

We have created a testbed version of the SAS, SAM, and CAS systems, with a small subset of the plates and imaging runs from the full survey. For testing development and changes in procedures, we use the testbed system.

13.4. Data Documentation

The details of the data in SAS and CAS will be documented. The Data Coordinator is responsible for assembling this documentation, with the assistance of the survey teams as necessary.

To describe the data at SAS, a data model is maintained in the “sas” SVN product and posted at http://sdss3.org/internal/datamodel. This data model contains a description of each archived file type and its required contents (and any known optional content), with the location of the file in the global directory structure, plus the meaning, data type and units of each entry. The data web server will have description of how to extract the data as well as examples of how to do so.

The algorithms used for producing reduced data will be documented on the http://sdss3.org site (as well as in technical papers of course).

To describe the data in CAS, each SQL table has descriptions, units, and data types for each item. In addition an algorithm page describes algorithms used in CAS itself. The CAS has documentation describing its use and as well as examples of how to do so.

In addition to providing online documentation, Jordan Raddick (JHU) will continue to offer programs to train the astronomical community in using the SAS, the CAS, and CasJobs to retrieve SDSS-III data. These programs will be based on the "Cooking with Sloan" trainings offered at previous AAS meetings, and the online follow-up materials for Cooking with Sloan will be adapted for SDSS-III. Training programs will be given twice a year as Special Sessions or Splinter Meetings at AAS meetings; there is also a small travel budget to pay for trainings at other astronomy meetings or invited trainings at university astronomy departments.

13.5. Data Releases

As observations progress, data archived on the SAS will be available to SDSS-III participants as it arrives, through the password-protected web server.

The CAS databases will be loaded periodically. Typically, we aim for the databases to be fully loaded six months or more previous to each public data release, to provide lead time for participants as well as an opportunity to vet the data. The Director and Project Scientist must approve the release of the data to the public.

The data release schedule is below. In most cases, we are aiming for a data release within a year after each observing season is completed. We include a DR11 "internal" release that includes data from the 2012-2013 season; with the full release scheduled for Dec. 2014, a public release the previous summer is unwarrant

Date	Data Release	APOGEE	BOSS	MARVELS[a]	SEGUE-2
Dec. 2010	**DR 8**	---	**Imaging** (up to Jan. 2010)	---	**Spectra** (up to July 2009)
July 2012	**DR 9**	---	**Spectra**[b] (up to July 2011)	**Radial velocities** (up to July 2011)	---
July 2013	**DR 10**	**Spectra**[c] (up to July 2012)	**Spectra** (up to July 2012)	---	---
July 2014	DR 11 [internal][d]	Spectra[c] (up to July 2013)	Spectra (up to July 2013)	---	---
Dec. 2014	**DR 12**	**Spectra** (complete)	**Spectra** (complete)	**Radial velocities** (complete)	---

Notes: (a) MARVELS will only execute a public release if additional funding is identified. (b) Initial BOSS spectroscopic reductions, released in DR9, will use the preliminary BOSS pipelines. (c) APOGEE intermediate data releases will only include completed integrations of stars (not partial integrations). (d) The internal data release DR11 is to help iron out problems and prepare the collaboration for the final release a few months later (this approach was invaluable for SDSS-II).

13.6. Data Distribution Reviews and Milestones

Here we lay out the various milestones for the data distribution effort, along with their planned completion time. The database loads are associated with acceptance reviews of data quality assurance tests and functionality. The testbed database is a small fraction of the total data that we use for experimentation and testing.

Archiving system for raw data developed: Summer 2008 (completed)

Archiving system for reductions developed: Fall 2008 (completed)

SAS web service available to collaboration: Spring 2009

Large-scale structure samples developed: Summer 2009

CAS database available to collaboration: Fall 2009

CAS mirroring tools completed: Spring 2010

Testbed database loaded [Acceptance review]: Summer 2009

DR8 database loaded [Acceptance review]: Winter 2010

DR9 database loaded [Acceptance review]: Fall 2011

DR10 database loaded [Acceptance review]: Winter 2013

DR11 internal release database loaded [Acceptance review]: Winter 2014

DR12 database loaded [Acceptance review]: Fall 2014

14. Quality Control

14.1. Introduction

Quality control of data are crucial, and in surveys of this size it must be automated, with quality expressed using well-understood metrics. We divide quality control into two categories: real-time assessment of raw data quality as an aspect of operations, and a higher latency determination of the quality of the final reductions.

14.2. Imaging Quality Control

The imaging QA begins on the mountain with the real-time data acquisition system. Several basic things are checked, including the bias levels and read noise of the CCDs, the FWHM of the images on all the CCDs, the sky levels on all the CCDs, and the form of the TDI flat fields. The read noise is checked against historical data and serious deviations are flagged. Biases are taken every imaging night and the levels and profiles are compared with historical data; again significant deviations are flagged. The sky levels vary a great deal, of course, but excessive levels caused by, for example, auroral activity, and misbehaving sensors can easily be found. Many of these checks are performed automatically and are flagged on a display maintained by a sophisticated piece of monitor software called Watcher.

There is a FWHM limit of 2 arcseconds for imaging, and it must be photometric. The decision to switch to imaging during a night is motivated by seeing measurements by the differential image motion monitor (DIMM), anecdotal evidence from the other telescopes on the site (primarily the 3.5-meter), the output of the 10-micron all-sky cloud camera, and, of course, by the observers' experience with the weather at the site. Often the observing will *begin* with imaging, motivated by the weather patterns. In any case, once an imaging run is underway, there is real-time feedback from real-time analysis of the image FWHM for stars automatically chosen from each CCD, and sensors with discrepant image sizes are flagged on Watcher. The Watcher display in this case is an array of boxes on the screen arrayed like the CCDs in the camera, with the sensor identification and the current FWHM displayed. The boxes are green if the FWHM is satisfactory, red if not; the green-red go/nogo theme is maintained for all of the displays. The transparency is evaluated in a rather indirect way by looking at the variance in the image obtained by the all-sky thermal infrared camera. This is calculated every five minutes, and there is a simple threshold set for satisfactorily photometric conditions; this has worked very well in practice.

Before processing, we determine a crude astrometric solution and store it. At the same time, basic imaging checks are performed on the status of focus, the rotator, and the scan rate (checking for suspicious "skippiness" between field centers); bad fields are flagged.

After photometric processing, the data are calibrated against overlapping existing data (including the rapid Apache Wheel cross-scans, which virtually guarantees overlaps). This "incremental" calibration allows us to confirm the photometricity decisions made by the observers on a short time scale. Eventually each run is included in a global recalibration self-consistently.

With photometric catalogs, we run a set of "runQA" checks for each camcol, which checks the position of the stellar locus, and the consistency of the PSF photometry with aperture determinations.

For each field we track its status through the photometric pipeline: whether it was flagged as bad in any of the checks; whether it was successfully processed, and if not why not; how it contributes to the final photometric footprint; how well the PSF was determined and whether any errors occurred; and how well it was calibrated.

14.3. SEGUE-2 Spectroscopy Quality Control

For real-time quality control, several checks are performed. First, observers perform detailed tests at the beginning of each run to determine the general health of the spectrographs, and they analyze calibration images for each plate for correct positioning, focus, and uniformity of focus.

Second, a "quick-look" spectroscopic reduction pipeline called SoS (Son of Spectro) is run at APO as soon as the data are obtained. Some specific tests for historical problems that SoS performs are: searching for excessive flux in the red LED part of the spectrum, and detecting problems in the serial data transfer. SoS compares the photometry to synthesized magnitudes from the spectroscopy, and measures gradients in this difference between the two. SoS provides real-time S/N evaluations at the end of each exposure. The observers use the S/N evaluations to decide when a particular plate is done. The observers also check the SoS results by eye by looking at subsamples of the resulting one-dimensional spectra.

Quality control on the final 2D reductions proceeds in a manner similar to that for the SDSS-II spectroscopic survey. Each plate is assigned a median S/N ratio based on the optimal extractions. The software automatically flags bad pixels and suspicious extraction features (many missing pixels, dichroic mismatches, etc). QA plots of throughput and other instrumental quantities as a function of CCD and focal plane position are checked by the reduction team for suspicious patterns; they then diagnose any errors that have occurred and flag the plates appropriately.

The SEGUE-2 team will define "survey quality" conditions that must be met for data to be considered complete and part of a homogeneous sample. The SSPP pipeline measures catalog uncertainties and flags suspicious cases. The reduction team checks QA plots of stellar parameters to evaluate the consistency of the data. They will compare multiple observations of the same star to check end-to-end consistency for different observing conditions.

SEGUE-2 team has defined "survey quality" conditions that must be met for data to be considered complete and part of a homogeneous sample. The primary requirement is that the final, delivered S/N for stars near the main sequence turnoff ($0.05<g-r<0.48$) near the faint magnitude limit, $19.25<g<19.5$, be at least 10. That is sufficient to get [Fe/H], log(g) and Teff for all stars within the Teff limits of the stellar parameters pipeline. The second requirement is that the zero point of the radial velocity for each plate, as checked by the second plate in each overlapping pair, be no larger than 7 km/s. That limit is set to be 3 sigma larger than the measured plate-to-plate zero point variation realized in SEGUE/SDSS-II measured using repeat observations in overlapping plate area as we plan to do in SEGUE-2.

Monitoring plates for delivered S/N is done by a combination of the PI, the observers and the data processing staff at Princeton. The requirements are set by the PI and the SEGUE-2 science team, implemented by the PI and data processing staff, and will be handed off over time to the observers.

Similar quality control mechanisms are used for the MARVELS survey pre-selection plates being taken concurrently.

14.4. BOSS Spectroscopy Quality Control

The uniformity of the BOSS spectroscopic survey will be ensured by integrating to desired S/N levels rather than for fixed integration times. This was made possible with the development of the near-real-time Son-of-Spectro (SoS) reduction pipeline during the first year of SDSS-I operations. This pipeline reports the $(S/N)^2$ at fiducial magnitudes for each exposure. Since $(S/N)^2$ scales linearly with exposure time for constant observing conditions, this allows the observers to tune the number of exposures and exposure times to achieve the desired S/N with an efficient use of telescope time. The fiducial S/N values and magnitudes will be determined by the BOSS team during commissioning.

Son-of-Spectro provides additional feedback for improving observing efficiency and identifying problems with the hardware, electronics or acquisition software. After each exposure, a plot is produced showing the relative throughput across the 3 degree field of view. This aids the observers in identifying guiding or plate scale problems, although the guider is the primary resource for this. Son-of-Spectro also checks for the re-appearance of any known problems to have occurred over the 8 years of observing. These include scattered light from light sources in the telescope or instrument LEDs, dropped pixels from electronics problems, defocusing of the spectrographs, and undue flexure of the CCDs relative to the instrument. The SoS pipeline is a subset of the full spectroscopic pipeline, and maintained by those developers. Any requested changes affecting SoS are tracked through the ticket system, approved by the Lead Observer, and developed with an explicit test for each change. New versions of SoS are installed at Apache Point by the observing staff during the monthly shake tests, falling back to previous versions if the test fails.

The calibration and health of the spectrographs are routinely checked. An afternoon checkout procedure tests the ability to take exposures and communicate with the spectrograph mechanicals. Bias, flat-field, and arc exposures are taken that measure the bias level, lamp brightness, and spectrograph focus. If need be, the spectrograph collimators are moved to move the position of the spectra on the CCDs or focus the instrument. A monthly checkout procedure takes a series of dithered flats for pixel-level flat-fields, observes a sparsely-plugged plate for characterizing scattered light, and observes a long integration on a blacked-out plate for characterizing dark currents. Changes to these procedures must be approved by the Lead Observer. The merit of any changes must be weighed against the additional burdens placed upon the observatory staff.

Quality control on the final spectroscopic reductions will proceed in a manner similar to that of SDSS and SDSS-II. As with the real-time reductions, a S/N value is measured at fiducial magnitudes in both the blue and red cameras. Our experience is that it is very rare for a plate to be declared done by the observers and require re-observing. The full reductions provide a more complete set of QA plots for diagnosing problems. The BOSS team will investigate any potential problems flagged either by the observers from the real-time reductions or from the full reductions.

The BOSS spectroscopy must ultimately deliver the redshift success rate for LRGs and Lyman-alpha forest noise properties as defined in the requirements documents. This will be tracked by the BOSS team through the life of the survey. Some objects will be observed multiple times on neighboring plates, which will be used for testing end-to-end consistency for different observing conditions.

14.5. MARVELS Spectroscopy Quality Control

During most of the 2008-2009 observing season, MARVELS quality control will be as described here. Stand-alone instrument interfaces and control software will be installed in January 2009 that will alter some of these procedures.

There is a MARVELS spectrograph checkout before each observing run. The observers clean the gang connectors, activate the MARVELS Java Control Interface (MJCI), and check its connection to the MC agents monitoring temperature, pressure, and PMT, and the agent controlling the calibration box. The observers check the temperature inside (such as fiber input mount, interferometer, grating, CCD house, bench center, and CCD area) and outside (such as room temperature and iodine cell temperature in the calibration box) the instrument.

In addition, there is a daily checkout. An IDL session started by the observer checks the status of the MC agents and the CCD. In particular, the following criteria for the CCD are checked: CCD temperature at -106 C, CCD pressure below 0.01, and that the CCD cooler has been activated.

During observing, real-time checks occur. Between each science image a bracketing exposure consisting of a Tio (tungsten lamp + iodine cell) and a ThAr lamp is taken. The observers check each image to verify that the brightest parts of the image are above 30,000 counts. Significant changes in the flux levels of these images can point to issues with flux transmission or aging lamps.

Observers check each image for missing spectra, which can indicate broken fibers or plate design difficulties. Quick-look software ("fluxqa") determines the signal to noise for each individual spectrum, recording also the LST, HA, seeing, RMS of seeing, and mean flux. During the 2008-2009 season, the relationship between RV error and spectrum signal-to-noise ratio will be calibrated and acceptance criteria developed by the MARVELS team. These acceptance criteria will be used to declare exposures done.

Until acceptance criteria are defined more rigorously, MARVELS requires an average S/N per pixel >20. In addition, the amount of flux accumulated in any acceptable observation must be >25% of the flux that would be accumulated in an observation taken under average observing conditions (i.e., we shall reject observations which are likely to suffer from error bars >2 times larger than average). Acceptable exposures of a field are marked with "OBSFLAG=OK" in the FITS image headers; others are marked with "OBSFLAG=JUNK".

Similarly, any spectra with average S/N per pixel <20 are flagged as dubious quality spectra in the output tables of RV's. An accumulation of more than three such flags on an exposure indicates a severe problem with broken or dropped fibers on that cartridge, and the pipeline earmarks those fiber numbers for inspection at the observatory and possible repair.

An uncertainty on each RV point will be calculated based on an error propagation resting on the theoretical error bars of the sinusoidal fits to the interferometer fringes in each spectrum- theoretical error bars which are calculated during construction of MARVELS whirls. These uncertainties will be used as a guideline in determining whether any observed RV variability is statistically significant. If the median RV scatter over all stars in a field is >2 times larger than the theoretical uncertainties, that may indicate the presence of unusual problems for that field, and we will earmark the field for a detailed check of the pipeline processing for that field.

14.6. APOGEE Spectroscopy Quality Control

For real-time quality control, the APOGEE team is developing the APOGEE Quality Assurance pipeline (AQuA). This software will calculate signal-to-noise ratios and other raw data quality checks for feedback to observers. The APOGEE team will develop additional tests of data quality. In addition, they are developing hardware check-out tests to be performed before observations. The software interface to the instrument will report instrument status.

APOGEE's final science reductions will calculate refined signal-to-noise ratio estimates. In addition, they will indicate overall quality values for plates as well as individual quality values for fibers. Bad pixels will be marked using quality masks. The APOGEE team will define "survey quality" conditions that must be met for data to be considered complete and part of a homogeneous sample.

14.7. Data Assembly Quality Control

The data assembly tasks mostly consist of matching among multiple samples (e.g. photometric and spectroscopic BOSS samples), and creating window functions and random catalogs for science.

Post-processing software in data assembly will measure the distribution of positional differences between matched objects across samples, as a function of imaging run and spectroscopic plate. Suspicious areas of difference will be flagged for eyeball checking. Physical parameters (flux, etc) that are common among different samples will be compared where appropriate, and outliers flagged.

Post-processing software for defining window functions will flag filled areas of the window function without corresponding science objects, and check for science objects outside the defined window function. For BOSS, early in the survey we will calculate benchmark luminosity functions and small-scale correlation functions; the post-processing software will recalculate these for each new sample to check for suspiciously large changes (those not attributable to improvement of statistics).

Estimates of total area, total number of runs and plates, and total numbers of objects will be calculated, published and sanity checked with other metrics of observing progress.

The period between completion of data assembly and the actual data release (typically six to nine months) is essential for vetting by the science teams and for addressing the resulting feedback.

14.8. CAS Quality Control

In the initial versions of CAS, we will verify that queries similar to ones run on the SDSS-II CAS can run at equivalent speeds and return similar results

For each version of CAS, automatic software will extract random sections of each table, and compare the results of each element numerically to those in the flat file distribution at SAS. Deviations larger than those attributable to round-off differences will be flagged for investigation.

Links among various data samples (e.g. photometry to spectroscopy) will be compared with the flat file distribution at SAS for consistency, again using automatic software.

The period between completion of the CAS and the actual data release (typically six months) is essential for vetting by the science teams and for addressing the resulting feedback.

15. Configuration Management and Change Control

15.1. Hardware Configuration Control

Responsibility for configuration management and change control during major hardware development efforts (i.e., BOSS hardware upgrades, APOGEE, MARVELS, and Infrastructure) resides with the respective team Principal Investigators and Instrument Scientists. CPO oversight, through the Change Control Board (Section 1.1.1), begins with the review, approval, and baselining of the survey SRDs. As hardware development moves through the design, construction, commissioning, and acceptance phases, the CPO and the Instrument Review Board will review the predicted and/or measured adherence of the system to the scientific performance requirements. If significant variances between performance and the requirements are determined, the CCB will convene to assess potential impacts and make recommendations of various programmatic actions for consideration by the Director.

In the course of conducting the surveys, a limited set of hardware improvement projects may need to be executed to bring performance into compliance with survey requirements. The work in each of these tasks is organized as a separate project with well-defined deliverables and schedules, and the lead responsibility for the completion of each project is assigned to a specific individual. Problems with existing systems are identified by members of the scientific and technical staff and filed in the SDSS-III problem reporting system. The issues are then discussed in the appropriate project teleconference and an individual is given an action to resolve it. If the resolution requires the definition of a new project, it is submitted to the Director, Project Scientist, and Program Manager for consideration and approval. If the new project is deemed necessary, the Program Manager works with the appropriate individuals to develop a cost and schedule estimate. Once the cost and schedule estimate is prepared, the project is reviewed against the current budget and schedule to consider if, when, and by whom the work will be done. New project requests with an estimated cost in excess of $3k are submitted to the Director for approval before they are incorporated into the WBS and work plan.

After new projects are approved, they are prioritized and assigned to an organization and individual, added to the WBS, and integrated into the project schedule. Institutional budgets are adjusted to fund a new project whenever the new project reflects a change in the scope of work that was defined in the original agreement between the institution and ARC. Funding for new projects must come from within the approved operating budget, either by postponing planned lower priority work, or by allocating some of the contingency reserve that is controlled by the Director.

15.2. Software Configuration Control

All software used in mission-critical production operations, from observing to data processing to data distribution, is under version control. All production software is stored in a source code control system and tagged with a unique version number. The source code control system is discussed in the next subsection.

The version of software used in production is under the control of the individual responsible for the project area in which the software is used. Production version upgrades are made in consultation with survey management.

We use the trac ticketing system (described in the next subsection) to identify and track approved software problems, changes, or improvements. Software bugs filed as Critical indicate a problem serious enough to negatively impact operations or adversely affect data quality. These problems are typically addressed promptly by the developer and implemented into production as quickly as possible. Other, less critical bugs are discussed during weekly project teleconferences and agreements reached regarding which bugs and/or change-requests will be worked on, and on what time scale. The decision to work on a particular bug or implement a particular improvement is in general made jointly by survey management, the developer, and the group who will receive and implement the modified code.

When a developer has completed the agreed-upon modifications, the code is checked into the code repository and the recipient operations group notified that the code is ready to test. The developer will typically provide the recipient group with instructions for testing and validating the improved code, and will often participate in the testing and analysis process. Once a new code version is successfully tested and validated, it is declared current by the production group and placed into production use.

15.3. Configuration Management Tools

We store all software in a source code control system, a "trac" system at https://trac.sdss3.org with an associated SVN repository. The trac system includes a wiki for documentation. SVN provides change tracking, concurrent editing, and version tagging of software products. Finally, we use the trac ticketing system for assigning tasks and reporting problems, categorized by survey component and relative priority. All survey participants are free to inspect software products and report problems.

The software packages in SDSS-III have significant interdependencies. We use the ExtUPS system to track these dependencies and to handle installed products. Each product has tables describing the other software it depends on. Standard "setup" scripts can set up a product to be used and automatically set up dependencies as well. The system can have multiple versions of the same product installed, as well as multiple builds for different operating systems of the same version.

A script "sdss3install" checks out from SVN, compiles, and installs into SDSS-III any valid SDSS-III product, allowing system adminstrators and users to easily install new software versions as they are developed.

With distributed data processing, special care must be taken to ensure version control is being obeyed. Experience has shown this to be a subtle problem: pipeline developers can easily be running multiple versions or special purposed code without realizing it, and sometimes only independent checks can reveal these problems. With this in mind, either the SAS or SAM institution will have all survey software installed, and will run spot-checks on the raw data to ensure that the software is being run as advertised. The Data Coordinator will be responsible for running these spot checks.

Legacy software from SDSS-I and -II will be converted and stored in the same repository, keeping the revision history. The Data Coordinator will be responsible for this conversion.

16. Safety, Environment, and Health

Policies for safety, environment, and health for SDSS-III are derived from the institutions where the work is performed, namely Apache Point Observatory (APO), Lawrence Berkeley Laboratory (LBL), University of Washington (UW), University of Virginia (UVa) , University of Arizona (UA), John Hopkins University (JHU), and University of Florida (UF). This work includes data acquisition, Plug Plate drilling, and instrument development.

New Mexico State University (NMSU) operates APO for the Astrophysical Research Consortium (ARC) and employs the permanent staff. The applicable NMSU policies are generally derived from, and accountable to, NM OSHA guidelines and regulations, and are documented at the following locations:

- http://www.nmsu.edu/~safety/
- http://www.apo.nmsu.edu/Site/usersguide/finalsafety.htm

The APO Site Operations Manager, is responsible to ARC for the safe conduct of all activities at the Observatory. The APO Safety Officer, provides safety oversight for all activities at APO, provides training and maintains the records. In order to fulfill this responsibility, APO provides the safety equipment and training for staff engaged in activities at the Observatory. To ensure the adequacy of the site safety program, external safety officers are called upon to perform periodic site audits.

Additionally, APO staff fall under the NMSU Environmental, Safety, and Health (EH&S) division within the university. NMSU EH&S review the safety plan and changes to it, review accidents and investigations, and perform annual safety audits of the facility. NMSU EH&S also provides many of the OSHA required trainings and disposal of hazardous waste. NMSU EH&S information is available at http://www.nmsu.edu/~safety/.

Some of the on-site visiting staff at APO are employees from various partner institutions. These individuals comply with both their institutions and NMSU/APO policies and observe the stricter of the two policies, in cases of conflict or non-overlap.

Instrument development will be accomplished at UVa, UF, UA, JHU, and LBL. Operations are conducted in accordance to their respective university safety guidelines.

- http://keats.admin.virginia.edu/
- http://www.ehs.ufl.edu/
- http://www.ehs.uci.edu/
- http://risk.arizona.edu/healthandsafety/index.shtml

Operations at LBL are conducted according to Department of Energy policies, which can be found online at the following locations:

- http://www-esh.fnal.gov/pls/default/esh_home_page.html
- http://www.fnal.gov/directorate/Policy_Manual.html#Section_3
- http://www-esh.fnal.gov/FIESHM_Plan/IESHM_011404.pdf

Plate-drilling operations are conducted according to the University of Washington Physics Department's policies, which can be found online at the following locations:

- http://www.phys.washington.edu/safety.htm
- http://www.ehs.washington.edu/

17. Education and Public Outreach

SDSS-III's Education and Public Outreach (EPO) effort will build on the success that SDSS-I and SDSS-II had in promoting science learning in both formal and informal education. In the realm of formal education, the SkyServer web site (http://skyserver.sdss.org) – the EPO interface to the CAS – offers the entire SDSS dataset, free of charge, to students and teachers, along with lesson plans for learning astronomy at the K-12 and introductory college level. In the realm of informal education, Galaxy Zoo (http://www.galaxyzoo.org) gives volunteers from the general public who are, in general, not scientists a chance to participate in a real astronomy research project, aiding to the scientific community, by classifying galaxies as either spiral or elliptical though a web-based interface. Other informal education resources, including links to image galleries, visualization tools, and the American Museum of Natural History's Science Bulletin feature on the SDSS, are linked to the SDSS web site http://www.sdss.org.

With the retirement of SDSS-II EPO Coordinator Julie Lutz, the EPO Coordinator of SDSS-III will be Jordan Raddick of Johns Hopkins University (JHU). The role of the SDSS-III EPO Coordinator will be to coordinate the EPO teams working on SDSS-III, to develop project plans and timetables, and to liaison with other facets of SDSS-III and with outside EPO teams. Lutz will remain available to SDSS-III EPO as a consultant.

We will continue to develop SkyServer as a resource for learning astronomy at the K-12 and college level. Raddick will continue to maintain the public data access tools on SkyServer; SDSS-III will also employ a programmer, based at JHU, to develop new public-friendly data access tools for SDSS-III data specifically. These new tools for viewing SDSS-III data will be released with the data releases that include those data.

Raddick will incorporate SDSS-III data into existing SkyServer lesson plans, and will create 1-2 SDSS-III-specific lesson plans for Data Release 8, as well as one additional project each for the new datasets offered by APOGEE and MARVELS (pending funding for a MARVELS public data release). In addition to creating these new tools and projects, Raddick will create a full-day professional development workshop to train K-12 college teachers and/or college astronomy instructors in how to use SkyServer projects to teach science. This workshop, which will be based on workshops that Raddick has previously led at meetings of the American Association of Physics Teachers (AAPT), will allow teachers to do authentic astronomy research using SkyServer tools, and will include follow-up through online learning tools (such as wikis) during the school year. The workshops will be created by January 2010; the first workshop will be held in summer 2010, with workshops held yearly until 2014. In addition to holding our own workshops for both audiences, we will apply for professional development credit for K-12 teachers with the state of Maryland, which would include funding from the state and assistance in finding interested teachers.

We will also continue to explore Galaxy Zoo's model of "citizen science" – online participation by volunteers from the general public who are typically not scientists in a real, publishable astronomy research project – as a means to promote general science learning in the populace. Galaxy Zoo is now evolving into a sustainable cyberinfrastructure program with the working title of "Zooniverse." Zooniverse will include data from the Hubble Legacy Archive and the Lunar Reconnaissance Orbiter. Zooniverse may not include SDSS-III data, but since Galaxy Zoo started as a program of the Sloan Digital Sky Survey, and the beauty of SDSS images was a motivating factor for Galaxy Zoo's 170,000 citizen scientists, there is a clear intellectual linkage between SDSS-III and Zooniverse. The goal of the SDSS-III citizen science efforts will be to conduct education research to learn about how citizen science impacts volunteers' understanding of scientific concepts and the process of science, as well as volutneers' attitudes toward science. This knowledge will help the scientific community design better citizen science projects in the future.

Zooniverse is being developed by a worldwide team led by Chris Lintott of Oxford University (UK). The education efforts of Zooniverse are centered at Adler Planetarium (the education director is Lucy Fortson), and coordinated by a team consisting of Fortson, Karen Carney (Adler), Nancy Dribin (Adler), Pamela Gay (Southern Illinois University at Edwardsville), Raddick, and Lintott.

Primary responsibility for the citizen science aspects of the SDSS-III EPO program will lie with the NASA-funded Center for Astronomy Education (CAE) at the University of Arizona, a partner in SDSS-III. The CAE members working on SDSS-III will be led by Ed Prather and Gina Brissenden. SDSS-III will fund the hiring of a postdoctoral researcher in astronomy education research at CAE (50% funded by SDSS-III for two years beginning in Fall 2009). The postdoc will be supervised jointly by Prather and Brissenden, with additional advising from Raddick.

In the first year (Fall 2009-Fall 2010) the postdoc will work with members of the CAE team, through its Collaboration of Astronomy Teaching Scholars (CATS) program. Members of CAE and the CATS program have already begun initial research into issues related to science literacy and attitudes toward science. The post will become an integral part of the CAE/CATS collaboration, and will begin conducting a systematic qualitative research program with Zooniverse volunteers to determine their state of knowledge and attitude about science. The results from this qualitative investigation will be used to inform the creation of the first generation of multiple choice questions. The outcome of this initial work will be used to inform the development of a instrument that can be used with citizen scientists pre- and post- their involvement with Zooniverse to measure the change in science literacy and attitudes toward science. He or she will then field-test this instrument with Zooniverse volunteers and make refinements as necessary. The final assessment instrument and results of this investigation will be published in the astronomy education research literature, and will be made available online to the community.

In the second year (Fall 2010-Fall 2011), the postdoc will use this newly developed assessment instrument to study whether volunteers undergo a positive change in understanding or beliefs as the result of participating in citizen science. He or she will recruit volunteers who are just beginning to participate in Zooniverse, and will give them the instrument as a pre-test. He or she will then follow the online posts of each participant on the forum, and will score the scientific understanding and attitude reflected in the posts. At the end of the year, the postdoc will give the test to the participants once again. As a measure of participant learning as a result of participating in citizen science, we will calculate the normalized gain from their pre and post test scores. These test results will be combined with the results from the forum study to develop a more complete picture of citizen science learning. Results from this study will be published in astronomy education research literature.

SDSS-III education work will be coordinated through phonecons held weekly while the postdoc is at Arizona during Fall 2009-Fall 2011, and biweekly at other times. The citizen science aspects of the program will be planned and updated at twice-yearly face-to-face meetings, either at the University of Arizona or at Adler Planetarium. Raddick, in his role as EPO director, will also attend SDSS-III collaboration meetings to update the rest of the larger project on EPO efforts, and to record and coordinate EPO efforts by SDSS-III scientists.

The position of Public Information Officer (PIO) established in SDSS-I will continue into SDSS-III; Raddick will also serve as PIO of SDSS-III. The duties of the PIO are to write press releases (in collaboration with the Spokesperson and the scientists whose work is highlighted), to coordinate the production of press releases and press conferences with PIOs at the respective participating institutions, and to stage press conferences at professional meetings. The PIO also helps with activities associated with AAS meetings (e.g. the SDSS exhibit booth), handles inquiries from the press, and provides an interface between the project and media (e.g. film crews at APO).

18. Risk Assessment and Management

Projects with the scope and complexity of the SDSS-III Survey must be carefully assessed to identify areas containing a high degree of risk. Risks may be associated with internal factors such as inadequate operations planning, equipment failure, complications associated with new development work that has never been done before, poor cost estimates, and the loss of key personnel. Risks may also be associated with external factors including weather and changes in economic conditions that impact the ability of participating institutions to meet financial and personnel resource commitments. The failure to identify and properly manage these risks may result in cost overruns, schedule slippages, and inadequate technical performance.

The key elements of a comprehensive risk assessment and management program include risk identification and analysis, implementation of risk minimization and contingency plans, and an appropriate level of contingency reserve. The following sections identify potentially serious risks by project area and discuss strategies for minimizing these risks. Contingency reserve and management is discussed in Chapter 19.

The quarterly reports from all survey units will contain updates on risk assessment, so that new items can be uncovered and old items can be tracked. On an annual basis, the Project Manager will bring together the members of the technical staff to review system performance during the previous year, identify areas of concern regarding the on-going condition of systems and equipment, and identify equipment or systems requiring improvement to maintain operational or science requirements

18.1. Survey Management

Survey Management comprises the activities of key project personnel, including the Director, Project Scientist, Program Manager, Spokesperson, and ARC Business Manager. The primary risk associated with Survey Management is the departure of one of these individuals, which could result in the loss of critical scientific and technical skills, and institutional knowledge. The risk is minimized by a steady dialogue and by the continuous sharing of information between these individuals.

Withdrawal of a financial commitment from an invested partner is unlikely because we have obtained signed Memoranda of Understanding with explicit pledges. Even if an institution were to withdraw, the impact would be small because of the large number of participating institutions. More serious (but still very unlikely) would be the withdrawal of one of the partners contributing essential operational resources. This risk is minimized by successfully meeting Survey milestones and goals and by maintaining good communication with each institutional partner to remain aware of changing conditions that may impact future involvement.

18.2. Observing Systems

Observing Systems includes all of the equipment and systems used to acquire data at APO. It also includes the personnel responsible for maintenance, support, plate plugging, and observing. A complete description can be found in Chapter 10.

Much of the equipment and systems associated with Observing Systems have been in place for several years and were used to acquire data for the SDSS and SDSS-II Survey. During the preparation of the SDSS-III work plan and cost estimate, we considered system performance and reliability over the past 8 years of operation, and the impact of aging and obsolescence, in our technical risk assessment.

Given the length of the overall SDSS project, with many key systems developed in the 1990's, loss of knowledge is a key risk. A number of long-time leaders of the technical side of the project are approaching retirement. SDSS-III must ensure transfer of knowledge to a new set of technical leaders *and* ensure that these new leaders do not themselves become single-point risks.

Operations and equipment at APO are covered under New Mexico State Risk Management. Buildings, and fire alarm monitoring systems along with boilers are inspected in-accordance with state risk management and state laws. All buildings and equipment are also covered by the state's insurance against loss or damage. This protects the project funding if damage occurs since much of the cost can be recovered.

18.2.1. Sloan Foundation 2.5m Telescope

The telescope is in good operating condition. Potential risks going forward include critical component failure (e.g., drive amplifiers, motors, bearings, encoders, mirror control and support systems, etc.), component obsolescence, damage to optics, and the loss of institutional knowledge.

A strong preventive maintenance (PM) program and a spare parts inventory system were put in place during SDSS-I operations, resulting in average system uptime of 96.5% over the past three years. The PM program and spare parts system will continue to be used for SDSS-III. Moreover, in preparing for SDSS-III operations, the spare parts inventory and documentation were reviewed. Additional components were purchased as necessary to provide additional sparing for six more years of operation. Documentation was updated, compiled, or created as necessary to provide sufficient information for continued operations. Protecting telescope optics is a very high priority. Procedures have been developed and vetted and are in place for handling telescope optics. In addition, all optics-handling operations are performed under the direct supervision of the Lead Engineer. We are reducing our schedule for aluminization of the primary mirror from every year to every two years in order to reduce the risk incurred by shipping the mirror to Kitt Peak National Observatory. To protect against the potential loss of institutional knowledge over time, we continue to cross-train site technical staff and ensure overlap and sufficient training when hiring replacement personnel whenever possible. We also utilize part time and emergency hiring options in order to help reduce and fill knowledge gaps.

18.2.2. Imaging Camera and Spectrographs

Both instruments are in good operating condition and should remain operational through the survey. Potential risks going forward include the loss of an instrument due to a lightning strike or other electrical surge, loss of a CCD, component obsolescence, and accidental damage during maintenance and repair.

The threat of lightning strikes or other electrical surges is mitigated through the use of conditioned power and isolation transformers that isolate the instruments from line power; comprehensive grounding connections for the telescope enclosure; and the use of either fiber optics or opto-isolator circuits to eliminate copper wire connections from penetrating the enclosure shell. The loss of a CCD in the imaging camera array would be problematic, as we have minimal spare CCDs and the time and effort to install, calibrate and properly test a replaced CCD would significantly impact the project schedule. Obsolescence concerns are addressed by maintaining suitable replacement components in the spare parts inventory at APO.

All work on the imaging and SDSS spectrographic cameras is done under the direction of Jim Gunn and/or Connie Rockosi. The cameras for the spectrographs and imager are continuously monitored for their power, cooling, and vacuum health. When critical items exceed pre-established values an email alert system notifies the staff as well as the instrument builders. If critical errors occur that can't be handled by site staff and would prevent the acquisition of survey-quality data, Jim

Gunn, Connie Rockosi, and technical support staff travel to the observatory to make emergency repairs. Less-critical problems are addressed on an annual basis during the scheduled summer shutdown.

We expect to retire the imaging camera after Fall 2009, when BOSS imaging should be complete. As of December 31, 2008, enough survey-quality BOSS imaging has been obtained such that a critical failure of the imaging camera would not significantly impact the BOSS survey.

The upgrade of the existing SDSS spectrographs to the BOSS spectrographs will greatly increase the pool of those with the expertise to repair them, and it reduces risk of obsolescence. To avoid the long downtime that would occur with the loss of a CCD, a spare science-grade device will be held for each of the blue and red cameras. This would reduce the risk of a catastrophic CCD failure to the several days of down-time required to install the spare CCDs in the existing dewars.

The MARVELS and APOGEE instruments are bench mounted and sit in rooms in the support building. They are designed to be highly stable and reliable and to minimize the amount of maintenance. Any maintenance inside the containers will be with the active supervision of the instrument teams. Procedures for the routine operations of the instrument will be documented by the instrument teams. We will maintain a spare part reserve of less expensive items. Detector failures would require a long replacement period. In the case of APOGEE, we could continue to operate with the remaining detectors; in the case of MARVELS, we would have to use our engineering grade CCD while procuring a new science grade one.

18.2.3. Cartridge Handling and Plug-Plate Measuring Equipment

Cartridge-handling equipment includes the hardware used to handle the spectrograph fiber cartridges in the APO plug-plate lab and move the cartridges to and from the telescope. Plug-plate measuring equipment includes systems such as the Fiber-Mapper, which maps the location of optical fibers in a specific plug-plate. All equipment and systems are in good operating condition. Risks going forward include premature equipment failure and component obsolescence.

A preventive maintenance plan is in place to maintain the condition of this equipment. The PM program allows us to monitor the condition of equipment over time and to take corrective action as necessary to replace and/or repair components showing signs of wear and premature aging. The risk of component obsolescence is minimized by maintaining a suitable supply of spare parts and a software repository. These systems will be substantially overhauled in the summer of 2009.

18.2.4. Plug-Plate Production Operations

Plug-plate operations include plate design at NYU and plate fabrication in the UW Physics Shop. Risks include the loss of institutional knowledge and the loss of machine tools and facilities at UW due to an environmental catastrophe (e.g., earthquake or building fire) or machine tool breakdown.

During operations, one individual at NYU is responsible for running the target selection code and generating the drill files used to fabricate plug plates. However, others in the data processing group know how to perform this operation; therefore, risk is minimized through cross-training and redundant knowledge. A short-term loss of production capabilities in the UW Physics Shop would most likely be caused by a problem with the large drilling machine used to fabricate plates, or in the large Coordinate Measuring Machine (CMM) used to make quality control measurements on finished plates. Based on past experience, it is not likely that either of these machines would take more than a few weeks to repair (the limitation being replacement parts availability), so we do not expect to lose more than one drilling run due to a serious machine breakdown. It should be noted that in the past eight years of operation for SDSS-I and SDSS-II, no drilling runs were lost due to a machine tool failure. Notwithstanding, the short-term loss of fabrication facilities is minimized by maintaining a sufficient backlog of plates that have already been drilled and shipped to APO for observing. We plan to maintain a three-month surplus of plates at APO. For MARVELS, it will not be possible to do this in the first year due to the need for additional pre-selection observations; however, after this initial selection phase, we will be able to define our plates well in advance. A catastrophic failure caused by environmental conditions would be addressed by re-locating plate fabrication to another University or a commercial machine shop. The limitation is finding a fabrication facility with a drilling machine large enough to handle the ~1-m diameter plug-plates.

18.3. Observatory Operations

Observatory Operations includes the observatory site infrastructure, the staffing and resources required to maintain the observatory at the level necessary to sustain operations. Potential risks include environmental conditions that could permanently damage facilities (e.g., forest fire, building fire, lightning strike, etc); poor management practices that violate agreements with federal and state agencies (e.g., activities and actions that violate conditions in the agreement with the Untied States Forest Service (USFS) for operating the observatory on Forest Service property); and poor management or work practices that result in serious personnel injury or environmental damage that result in a shutdown of site operations.

To minimize the risk associated with forest fires, the Site Operations Manager has on several occasions obtained permits from the USFS and cleared timber away from observatory buildings and telescope enclosures. Increasing the distance between the forest and site structures provides a modest buffer zone around the site. In addition, the USFS occasionally thins areas near the observatory to eliminate dead-fall and remove diseased trees and excess brush. Additionally fire protection systems have been incorporated as well as building designs to reduce the possibility and extent of damage caused by a forest fire.

To minimize personnel risks associated with weather, the Site Operations Manager has implemented procedures that require personnel to remain indoors during severe weather. In addition, visitor guides available at the observatory warn visitors to take cover in the event of severe weather.

APO has in place procedures and plans in accordance with NMSU's Emergency Action Plan. This sets up authority and responsibility that addresses any emergency threat at the observatory.

To minimize lightning risk to site infrastructure, a lightning protection upgrade was completed in 2000. The upgrade was in response to recommendations from an external review committee charged by the Site Operations Manager with reviewing the status of lightning protection measures in place at the time. The upgrade helped ensure that electrical systems were properly isolated and grounded. Follow-up work included replacing the site telephone and fire alarm system to eliminate copper wires from running into sensitive areas; installing ground cables on all buildings and telescope enclosures; improving the grounding of the 2.5m telescope; and subscribing to an early-lightning-detection and warning system.

The Special Use Permit currently in force with the USFS goes to April 2018. To maintain a good working relationship with the USFS and preclude complaints, the Site Operations Manager works closely with the USFS to ensure that any and all site work is performed in accordance with USFS requirements and guidelines, and that all necessary work permits are obtained before work commences.

To minimize risks associated with poor management or work practices, all work is performed at APO in conjunction with the APO Site Safety Plan. The plan defines personnel responsibilities and accountabilities; identifies applicable codes, agencies, and inspection procedures; and outlines requirements and procedures for performing work in a safe and responsible manner. Front-line supervisors are responsible for ensuring that all work done under their direction is performed in a safe manner and in accordance with the Site Safety Plan. Additional information is provided in Chapter 16.

Another risk with Observatory Operations is the loss of key personnel, in particular a loss of observing staff such that we are not able to run the telescope full time. The key mitigation step is to maintain an observing staff of adequate size, so that stress does not need to loss of morale and so that we can maintain adequate staffing even in the face of illness, unplanned departures for other jobs, and so forth. We have long experience from SDSS-I and II of the size of the observing staff that is required, and we have drawn on this experience for our staffing plan.

18.4. Development Projects

New development work associated with the SDSS-III Survey is discussed in Chapters 4-9. Risk is typically higher on new development work, as there is greater uncertainty in the actual work to be done and amount of effort required. To minimize these risks, the scope of work and cost estimate associated with each area of new development was prepared in consultation with the PI and unit leads, individuals who have done similar work in the past, and the individuals who would be performing the work. A WBS was prepared for each development project to outline major deliverables and specify individual tasks at a reasonable level of detail. The WBS was then used to develop a cost estimate and budget for the work, and a schedule against which to track performance. The WBS for new development work is discussed in Chapter 3.

The Program Manager is responsible for developing and maintaining the WBS for new development work, in consultation with the PIs and unit leads. He is also responsible for tracking performance against the work plan and for working with the PIs and unit leads to make adjustments as necessary to ensure that development work is completed on time and within the approved budget.

All of the SDSS-III development projects must place a high value on performance, schedule, and risk management. The primary mechanism to achieve these virtues is the large body of experts within the SDSS-III collaboration coupled with the project management structure. Instrument teams are held to reviews, as described in Chapter 4. The Instrument and Software Boards provide a mechanism for hardware and software experts from across the SDSS-III collaboration to ensure a commonality of practices and to recognize problems before they become serious. In addition, it should be noted that most of the development projects are upgrades of existing systems where our team retains the experience from past builds.

We next discuss some of the more significant risks in each of the builds.

18.4.1. BOSS

The schedule for the BOSS spectrograph upgrade is important to hold. BOSS must be fully commissioned by January 2010 so as not to lose any opportunity to observe the North Galactic Cap. The schedule calls for the spectrograph upgrade to be completed in August 2009, with commissioning in September 2009. It is possible that on-sky testing of the performance of the instrument or of the target selection algorithms could require an alteration of target selection. It would require a month of observing to assess this and another month to alter the target selection and drill new plates. Hence, we must be on sky in November 2009. As such, the schedule has two months of contingency.

The dominant risk in holding this schedule is the procurements of the CCDs, the VPH gratings, and the replacement fibers. The blue CCDs are scheduled for delivery from e2v in April and May of 2009, with an engineering device in December 2008.

The red CCDs are to be provided as an in-kind contribution from LBNL. These CCDs are a new 4096x4096 format for LBNL, and entail the more serious modifications to the SDSS electronics by reversing the polarity and adding an 80V substrate voltage. We addressed this risk by fabricating a prototype dewar to house an engineering-grade LBNL CCD to test the CCD packaging, electronics and read noise. This test was successfully completed in December 2008. The three science grade LBNL CCDs (including 1 spare) are in fabrication and will be packaged by February 2009. The BOSS Instrument Scientist is PI of the LBNL laboratory that is producing the devices, so the lines of communication are excellent.

The red VPH grisms were due to be delivered in June 2008 but have not yet been made, due to delays in the vendor commissioning their larger bench. We are in frequent conversations with the vendor and have visited. The blue VPH grisms will be ordered as soon as the red ones are complete.

The fiber upgrade occurs in 8 units for the 8 "dark time" spectroscopic cartridges, to be delivered starting in March 2009. We do not view this as a major schedule risk, as problems will be uncovered early and we could begin operations with fewer than 8 cartridges. If the delivery were late enough or the performance poor enough to force a reassesment of the procurement, then we would face a schedule risk. We addressed this risk by prototyping these fibers with 3 vendors, and selecting the vendor with the best performance. The prototype fiber bundles were delivered on schedule and underwent extensive testing at the University of Washington.

The data acquisition system for BOSS is still in the design stage but is a straight-forward application of off-the-shelf parts and minimal recording software. If there are delays, the new BOSS CCD system has been demonstrated to work with the present SDSS-II data acquisition system; this is a viable fallback poisition.

The rest of the upgrade is a straight-forward extension of the previous instrument. We will minimize risk here by keeping these builds off the critical path.

The BOSS data reduction pipeline must have a high-level of performance to characterize and control systematic errors for the Lyman-alpha forest clustering analysis. The primary risk mitigation is to ensure that strong scientists are hired for the pipeline work. The project needs to enlarge the group of scientists that understand the inner details of the pipeline and ensure that this group communicates frequently to assess the ongoing pipeline work. These scientists must communicate well with the Lyman-alpha forest working group to address those stringent needs.

18.4.2. APOGEE

APOGEE has significant cost and schedule risks as it is still in the design phase and must be ready for survey operations by April 2011. Because SDSS-III ends observations in June 2014, delay in beginning the survey impacts the data acquisition substantially. APOGEE is particularly interested in the Galactic bulge, which is best observed in early summer. The APOGEE schedule is constructed to permit 4 seasons observing the bulge, but delays would threaten this.

The APOGEE instrument will consist of procurements for a large number of subsystems, which implies risk in coordinating many different contracts and interfaces. We will employ a half-time project manager to handle these procurements and track the resulting schedule. We will require clear documentation of all interfaces.

The APOGEE camera is a large optical system (40 cm optics). It will be built by an industrial partner on a fixed-price contract. We have been quoted delivery times of 9 months. This is acceptable, but we will meet frequently with the vendor to ensure that problems are detected early. We also will consider ordering the optical blanks prior to CDR so as to accelerate the schedule and purchasing a spare blank to mitigate schedule risk in the event of damage to an optical element during polishing.

The VPH grism is large and will require mosaic techniques that have not been demonstrated yet. We are working with 2 vendors to design the grism.

The APOGEE detectors are state-of-the-art Teledyne HgCdTe devices. Fortunately, we will receive a loan of 3 detectors from the University of Arizona. This avoids a significant schedule risk on a major procurement.

Successful design and integration of the APOGEE instrument depends heavily on key personnel, with the risk that these personnel are overloaded. We need to complete our hiring plan so as to spread the workload and hold schedule. Having an overloaded group risks shortcuts being taken as regards documentation and management practices, which increases the risk of errors.

The APOGEE data reduction pipeline includes an ambitious automated comparison to stellar spectral synthesis models. We intend to hire our personnel early to begin work. We also will cultivate the involvement of all of the spectral synthesis experts in the collaboration in the oversight of the pipeline development.

18.4.3. MARVELS

MARVELS hardware development is largely complete as of December 2008. However, work remains to put the survey into a smooth operating state and to achieve the desired level of the stability in the radial velocities produced by the instrument and the data pipeline.

The long-term stability of the instrument would be compromised by major failures that cause the instrument to be opened and repaired. This robustness was an important aspect of the design, e.g. to minimize moving parts. We intend to keep a close connection between the APO staff and the MARVELS instrument group, centered on the weekly operations phone conference.

MARVELS requires frequent handling of its fibers, particularly due to the use of a fiber coupler. Breakage of fibers is a risk that we must mitigate with carefully documented handling and maintenance procedures.

The long-term success of the MARVELS project depends strongly on the continued participation and increased leadership of a number of faculty and junior personnel, particularly as the survey moves from its development phase to an operations and science phase. We need to continue to cultivate an environment for the junior people to move toward leadership roles and for the broader SDSS-III collaboration to become involved. Early science results are an obvious way to continue the enthusiasm for the project.

18.4.4. SEGUE-2

SEGUE-2 development is largely complete as of December 2008. The largest risk is retaining key personnel to process the data and prepare the data release. Maintaining the collaboration communication and completing the hiring plan is the primary mitigation mechanism.

18.4.5. Infrastructure

The upgrades to the cartridge and fiber infrastructure are being conducted by the highly experienced group that built the original SDSS system. Nevertheless, there are a number of risk items.

The fiber system upgrade is a substantial procurement and carries a performance risk. We conducted a competitive bid between three vendors, including prototype units, in order to minimize the risk of low performance. We will also be testing the units as they arrive to diagnose problems early.

The other aspects of the upgrade comprise many pieces and the dominant risk is the management task of the distributed effort possibly leading to schedule slip. We are maintaining a resource-loaded schedule to address this risk.

The operations software controls the telescope and instruments and integrates data from various monitoring systems. It is being rewritten into a more modern language (Python) to accommodate the larger variety of instruments and observing strategies in SDSS-III and to make it easier to maintain in the future. The new system is heavily modeled after the current APO 3.5-meter telescope system and is largely being written by the team that wrote that software, so the technical risks are small. The group is holding weekly phone conferences to ensure coordination and to keep focus on their delivery schedule. If there are severe schedule slips, the BOSS system can be run with the existing software with minimal changes and MARVELS can continue to run in the stand-alone manner that it is using in the first year.

18.5. Data Processing

Data processing includes the activities and systems associated with processing and calibrating data for the four surveys. Potential risks include outputting substandard data due to inadequate quality control measures; schedule delays due to inadequate configuration control, data re-processing, and equipment failures; and loss of institutional knowledge.

Data outputs from the four surveys must be carefully inspected to ensure that the data meet the scientific goals for each survey, as failure to promptly and properly assess and verify data quality may have serious cost and schedule impacts. To minimize this risk, quality control measures are implemented at the observatory and at various steps in the data processing process, as discussed in Chapter 14.

Configuration control is necessary to maintain data consistency and maximize the use of data processing resources. In order to meet schedule requirements, data must be processed and calibrated in a production manner. This requires limiting changes to only those required to meet the approved science objectives and achieve the level of efficiency necessary to meet schedule requirements. Code changes and improvements require testing and validation, which consume resources. Moreover, if a change is of significant consequence, proper implementation may require the re-processing of existing data. Implementing and validating code changes, and re-processing data, consume resources and have the potential to impact project schedule and cost. To minimize these risks, configuration control measures have been put into place, as discussed in Chapter 15.

To minimize schedule risk due to equipment failures, we have built redundancy into the science archive by providing a full mirror at NYU. We also backup to the HPSS tape system at NERSC. The data processing itself is farmed to arrays of computers; if one computer fails, data processing will continue, albeit at a slightly slower pace until the faulty machine is replaced. Data are stored on hard drives configured in Redundant Arrays of Inexpensive Disks, or RAID arrays. Through the use of redundancy, RAID arrays offer increased data security, improved fault tolerance, and improved data availability. We protect our source code and document repository and wiki with a daily off-site backup, in addition to the NYU mirror.

To minimize risks associated with the loss of institutional knowledge, we cross-train personnel and update system documentation as processes change. We also work to automate many of the data processing operations, which in addition to improving operating efficiency, serves as a form of process documentation.

18.6. Data Distribution

Data distribution includes the activities and systems associated with preparing and loading processed data onto fileservers and into databases, and serving the data to the collaboration and general public. Potential risks include loading and distributing faulty data; failure to meet data release deadlines due to inadequate configuration control; failure to meet release deadlines and system availability requirements due to equipment failures; and loss of institutional knowledge.

As with data processing, sound quality control processes are required to ensure data integrity prior to and after loading the data into databases and onto file servers in preparation for data distribution. Loading faulty data means wasted time, as the data have to be reloaded once problems are discovered downstream in the loading and preparation process. Distributing faulty data is bad for any number of reasons. To minimize these risks, quality control processes will be built into the production data distribution operations. In addition, we make each data release available to the collaboration a few months before it is made available to the general public. This provides a period in which collaboration members use the data and identify problems. Depending on the nature of a given problem, we either fix the problem in the data or document its existence in the data release documentation. The evaluation period helps ensure the quality of released data products.

In designing the SDSS-III SAS and CAS systems, we have imposed a greater parallelism between the flat file distribution and the CAS SQL tables. In the database loading processing, new calculations are not allowed: all loaded quantities must also exist in the flat file distribution. This means that the flat files will be a complete representation of the data in the database, in contrast to SDSS-I and II in which the database did extensive calculations during loading to establish the links. We expect that the new system will be substantially easier to develop, maintain, and validate, and hence a notable reduction in risk.

Configuration control is necessary to minimize schedule risks associated with mismatched data models. When changes are made to data processing code and the data model, and those changes are not properly propagated into the data distribution operation, problems inevitably occur during the data loading process. To minimize schedule risk, configuration control measures to manage data model and database schema changes have been implemented.

To minimize risks associated with equipment failures, we have built redundancy into our data distribution systems. All fileservers are configured with RAID arrays. Multiple copies of each data release are spinning on different machines. In the case of the CAS and SkyServer, to avoid a catastrophic failure at a single site, we will have mirror copies in Europe, Asia, and South America in addition to the original JHU version.

To minimize the loss of institutional knowledge, we cross-train individuals, automate processes, and strive to keep process documentation updated.

19. Contingency Management

We have budgeted SDSS-III with a Management Reserve of approximately 10% of the original project budget. This value was computed by assigning different levels of cost risk to the different budget units: 5% for central management, APO operations, and MARVELS; 10% for SEGUE-2; 12.5% for Data Distribution; 15% for BOSS and the infrastructure development; and 20% for APOGEE. We are not holding reserve for the E/PO effort. Contingency is not rigidly tied to these project units but is available for the Director to assign where most needed. The Director will report on the distribution of the management reserve in the quarterly reports.

If additional cash became available, it could be allocated to pursue new projects (New Projects Fund), an example of one such possibility being support for designated tasks associated with Education and Public Outreach. If at any time the project suffers an emergency where the Management Reserve proves to be insufficient, then the remaining cash in the New Projects Fund can be tapped for this purpose. Conversely, if in later years it appears that holding a smaller Management Reserve is acceptable, then cash can be moved from the Management Reserve to the New Projects Fund. In any case, we do not intend to start new projects with major expenditures early in the project, to ensure sufficient balance to address potential large unforeseen expenses.

Prior to the submittal of each year's budget request to the Advisory Council, the Management Committee will assess any likely demands in the coming year on spending from the Management Reserve, based on an evaluation of the status of the project infrastructure and any other issues that are relevant. This assessment forms the basis for the proposed level of the Management Reserve that is part of the annual budget request.

Separately, the Management Committee will also consider the merits and priorities for activities that could be funded from the New Projects Fund. Based on that review, it makes a recommendation to the Advisory Council for specific expenditures. This recommendation also appears in the budget request, and requires the approval of the Advisory Council because any such funded activity represents an extension of the scope of work for the SDSS-III.

The above process ensures that the Advisory Council can monitor the level of the contingency, and can take an active role in the design and approval of extensions to the scope of work. It also provides for flexibility in the best allocation of resources as the survey progresses.

20. Program and Business Management System and Controls

20.1. Introduction

Comprised of the elements described below, the SDSS-III Program and Business Management System consists of defining documents, procedures, software tools, and manpower. These will be integrated and cross-connected so as to provide programmatic and business management reporting and planning support to the SDSS-III Director in the areas of survey operations and data management, technical development projects, schedules, budgets, cost accounting, contracts, and risk management.

20.2. Project Schedules

We have constructed detailed schedules for the development projects, survey operations, and data processing. These schedules are derived from work plans and estimates provided by team leaders and will be linked to elements in the work breakdown structure. They are rolled up into a survey overview schedule. The schedules are constructed using Microsoft Project (or the compatible Mac-based application, OmniPlan). The schedules will be cost and resource-loaded. They show dependencies so that critical-path method management can be utilized for project tracking and earned-value assessment.

Schedule slips are a key risk for SDSS-III not only because of the delay in or loss of survey science but also because of the salary overruns that they imply. Frequent monitoring of the schedule is our principal mechanism for early detection of such problems. All teams will report their progress relative to the schedule in their quarterly reports. Higher priority tasks, such as critical path items, will be monitored through the weekly MC phone conference. The Program Manager will communicate with all development projects monthly in regards to the schedule. The schedule will be updated quarterly by the Program Manager, Project Controls Specialist, and Project Managers based on progress reports furnished by the teams.

20.3. Planning Budgets

The survey planning budgets exist in Excel format, along with a document that gives the justification for each budget item. These documents are in Appendix C. The planning budget lists costs on a quarterly basis so that annual sums can be performed by fiscal year, calendar year, and academic year as needed. Inflation of 3% annual has been assumed.

Each row in the planning budget carries a unique tag that allows it to be cross-linked to the justification document. All survey expenses will be assigned to one of these tags, and the costs will be accumulated quarterly. The budget sheets contain separate columns for the planned and actual costs, so that comparisons can easily be made. The rows will also be linked to SSP codes (see section 20.4) as appropriate. While the planning budget is not as fine-grained as the WBS, WBS labels can be attached to budget rows.

The planning budget spreadsheet is actively linked that agency-specific budgets can be constructed and suitable what-if scenarios investigated.

20.4. MOUs and SSPs

The signed institutional MOUs are kept at ARC under the curation of the ARC Business Manager; the Director and Program Manager will keep copies. These are generally unchanged for the duration of the survey, as they define the general terms and conditions of survey participation, payment terms, and provisions for in-kind contributions, if any.

Sloan Survey Project (SSP) agreements, also called Support Agreements, are contracts between ARC and the supporting institutions that define work packages that the institutions provide to the project in return for payment or in-kind credit. They are executed on an annual calendar-year basis, and contain a statement of work, schedule, budget, reporting requirements, and a designated responsible manager at the SSP institution. The deliverables defined in the SSPs are linked to the planning budgets and to the accounting system. We anticipate approximately 20 SSP agreements.

The SSP agreements stipulate that the Director will pre-approve individual capital expenditures above $3k as well as all in-kind contributions. In-kind contributions will consist of capital items or effort that correspond to items in the planning budget.

20.5. Quarterly Technical and Financial Reports

The four survey PIs, the Data Coordinator, the Survey Coordinator, the Infrastructure Lead, the APO Site Operations Manager, and the Spokesperson file quarterly reports to the Central Project Office regarding the technical progress on their portion of the project. This includes milestones and goals met, problems encountered, science successes, survey progress, as well as risk assessments and goals for the coming quarter and updates to the project schedule.

The SSP managers provide quarterly technical progress and expenditures reports, in text and Excel format, respectively. The reports encompass progress and expenditures in the past quarter, with plans and projections for the present quarter. These are submitted to the Program Manager and Business Manager.

The Director, Program Manager, and Business Manager assemble these reports into a Survey Quarterly Report for submittal to the funding entities and ARC officers.

20.6. Cost Accounting Between Surveys

In accord with instructions from the NSF, project costs will be ascribed to the four individual components (BOSS, APOGEE, SEGUE-2, and MARVELS) according to the following flow-down. If a cost would be required for a SDSS-III scope that included only BOSS, then the cost will be ascribed to BOSS. This includes the understanding that in many cases one cannot efficiently hire a fraction of a FTE to do certain tasks. Continuing, if the cost would not be required for a BOSS-only project scope but would be required for a BOSS plus APOGEE project scope, then the cost will be ascribed to APOGEE. If the cost is required for BOSS plus APOGEE plus SEGUE-2, then it will be ascribed to SEGUE-2. If the cost would not be required for BOSS plus APOGEE plus SEGUE-2, but is required for MARVELS, it will be ascribed to MARVELS. It is recognized that this incremental approach to project costs will yield a substantially different accounting than would be found by splitting costs proportionally by telescope usage.

To insure that the SDSS-III cost accounting is accurate and that the agency funds are only used to pay for pre-approved expenses the support agreements will be written such that the above flow-down accounting policy is manageable. For some SDSS-III supporting institutions, this will require multiple support agreements, one for each SDSS-III project being supported: BOSS, APOGEE, SEGUE2 and MARVELS. For institutions supporting multiple SDSS-III projects, where the detailed costs can be categorized by project by more detailed accounting means, this may require only one support agreement but evidence of the required detailed accounting capabilities must be demonstrated in advance of a SDSS-III single support agreement being set-up. ARC has separate accounting for its other 3.5-meter telescope project as well as the SDSS-II archive project; there is no comingling of SDSS-III funds with other ARC project funds. NSF funds are received into a separate NSF-funds-only checking account so the expenses paid with those funds can be easily distinguished from expenses paid with SDSS-III member and other sponsor funds.

20.7. Revenue/Expense Reports

Verification of the Revenue and Expenditure Reports is a critical control for fiscal management to ensure that revenue and expenditure transactions are correct, allowable, and applied to the appropriate accounts. Review includes a reconciliation of expenditures comparing the source document to the report entries. The Summary Rev/Exp Report will be produced and distributed to the Advisory Council and Management Committee on a quarterly basis, monthly if requested. The three elements of the Summary Rev/Exp Report include:

- Revenue by Source (funding agencies and member institutions)
- Expenses sorted and summarized by support agreement, work breakdown structure summary codes, and planning budget tags and
- Cash Balance and Distribution of Cash Balances.

The Detailed Rev/Exp Report data, entered at the invoice payment level of detail, will be available to the Management Committee on an as-required basis. The Business Manger reviews invoices for accuracy and pays the invoices from the appropriate funding source. For example, only expenses included in the approved NSF budget plan are paid for from NSF funds.

The planning budget will be updated quarterly to report the expenses incurred for each budget item. This will allow a detailed comparison of costs relative to the planning budget.

20.8. Staffing

Staffing of the program and business management functions consists of the Program Manager (full-time), Business Manager (67%), Project Controls Specialist (full-time), and Development Project Managers (~2 FTEs budgeted during development time frames). Within the survey teams, there are experienced project leads that will support and contribute to the management system. Members of the Central Project Office, especially the Director, also contribute to management functions.

21. Appendix A: Baseline Survey Metrics

Figure A.1 presents plots of the baseline forecasts for the four spectroscopic surveys corresponding to the metrics discussed in Section 2.6. These forecasts are based on the following assumptions:

1. The division between dark time and bright time surveys is as specified in Section 2.5. As noted in Section 2.5, BOSS spectroscopy is given grey time when the North Galactic Cap is observable and dark time only at other times.
2. Weather is acceptable for spectroscopic observing 50% of the time, except for July when it is 40%.
3. The telescope is shut down in August for maintenance.
4. One bright night per month is spent on engineering.
5. SEGUE-2 exposures are 3.2 hours, including overhead. In Fall 2008, 30% of the dark time that would otherwise be available to SEGUE-2 goes to BOSS imaging.
6. BOSS spectroscopic exposures are 1.7 hours including overhead.
7. MARVELS spectroscopic exposures are 1.2 hours including overhead.
8. APOGEE spectroscopic exposures are 4.5 hours (cumulative) including overhead.
9. APOGEE plates have 250 fibers devoted to program stars (counted in the metric) and 50 fibers devoted to sky or calibration sources.
10. After APOGEE starts, MARVELS has an efficiency equivalent to using 75% of the available observing time. This is an approximation that will have to be sharpened once a detailed APOGEE/MARVELS co-observing plan is developed; note that our performance metric does not yet include the "quick-look" survey that MARVELS will carry out in the APOGEE "short fields" (using the remaining 20-25% of the time).
11. APOGEE observes during 95% of the allocated bright/grey time, with the remaining 5% available to MARVELS fields that may not be desirable for APOGEE.
12. SEGUE-2 runs from September 2008 through July 2009. BOSS spectroscopy (following commissioning of the upgraded spectrographs) runs from December 2009 through June 2014. MARVELS runs from September 2009 through July 2014. APOGEE runs from April 2011 through July 2014.

Note that SEGUE-2 plates have 640 fibers, BOSS plates have 1000 fibers, and APOGEE plates have 300 fibers. For MARVELS we assume 60 fibers for the first two years and 120 fibers for the two subsequent two-year cycles. For detailed explanation of the MARVELS and APOGEE metrics, see Section 2.6.

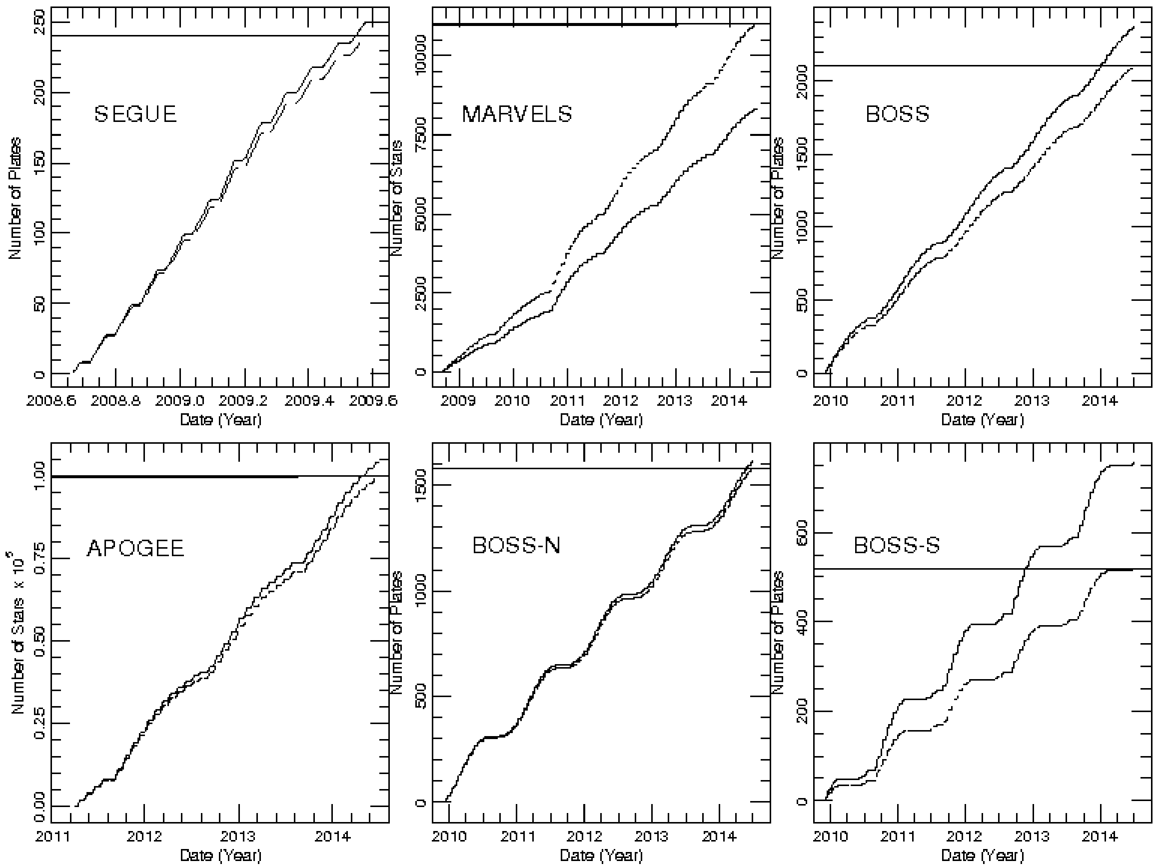

Figure A.1: Forecast of performance relative to metrics described in Section 2.6.

In each panel, the solid curve shows the forecast of data acquisition based on the assumptions above. The horizontal solid line shows the survey goal. The dashed curves are the forecasts scaled by a constant factor so that they achieve the goal at the end of the survey: if we keep pace with the dashed curve, the survey will achieve its goal. For BOSS, we separately show results for the northern and southern survey area. If the amount of BOSS southern imaging exceeds the original goal, then some northern area will likely be traded for southern area, which would lower the horizontal line in BOSS-N, raise it in BOSS-S, and give us more margin for achieving the total desired area.

22. Appendix B: Work Breakdown Structure

This appendix contains the detailed Work Breakdown Structure of SDSS-III. The top-level WBS is given in section 3.

23. Appendix C: Planning Budget and Justification

This appendix contains the planning budget spreadsheet and the corresponding justification. The planning budget consists of ten budget sheets, each developed from a bottom-up costing of portions of the project following the organization chart.

24. Appendix D: Project Schedule

This appendix contains the detailed project schedule files.