SDSS-III: Massive Spectroscopic Surveys of the Distant Universe, the Milky Way Galaxy, and Extra-Solar Planetary Systems[1]

January 8, 2008

[1]Abstracted from a proposal to the National Science Foundation, submitted in September, 2007

EXECUTIVE SUMMARY

Building on the extraordinary legacy of the Sloan Digital Sky Survey (SDSS) and SDSS-II, this document presents a six-year program (SDSS-III; mid-2008 to mid-2014) that will use the wide-field 2.5m telescope at Apache Point Observatory to carry out four surveys on three scientific themes: dark energy and cosmological parameters; the structure, dynamics, and chemical evolution of the Milky Way; and the architecture of planetary systems. The Baryon Oscillation Spectroscopic Survey (BOSS) will measure redshifts of 1.5 million luminous red galaxies and Lyman-α absorption towards 160,000 high redshift quasars. By using the baryon acoustic oscillation scale as a physically calibrated ruler, BOSS will determine the absolute cosmic distance scale with precision of 1.0% at $z = 0.35$, 1.1% at $z = 0.6$, and 1.5% at $z = 2.5$, achieving tight constraints on the equation of state of dark energy. The high-precision clustering measurements over a wide range of redshifts and length scales will also provide rich insights into the origin of cosmic structure and the matter contents of the universe. SEGUE-2 will use the SDSS spectrographs to measure radial velocities, spectral types, and elemental abundances of 350,000 stars in numerous target categories to a magnitude limit $g \approx 19$, probing the kinematics and chemical evolution of the outer Milky Way. The APO Galactic Evolution Experiment (APOGEE) will use high-resolution ($R \sim 20,000$ at $\lambda \sim 1.6\mu$m), high signal-to-noise ($S/N \sim 100$) infrared spectroscopy to penetrate the dust that obscures the inner Galaxy from our view, measuring radial velocities, spectral types, and detailed elemental abundances of 100,000 red giant stars to a magnitude limit $H \approx 13.5$, across the full range of the Galactic bulge, bar, and disk. Together, SEGUE-2 and APOGEE will provide a picture of the Milky Way that is unprecedented in scope, richness, and detail; the combined data set will play a central role in "near-field cosmology" tests of galaxy formation physics and the small scale distribution of dark matter. The Multi-Object APO Radial Velocity Exoplanet Large-area Survey (MARVELS) will use fiber-fed interferometric spectrographs to monitor the radial velocities of 11,000 bright stars, with the precision and cadence needed to detect gas giant planets with orbital periods ranging from several hours to two years. With a unique combination of enormous numbers and well characterized sensitivity, MARVELS will provide a critical statistical data set for testing theories of the formation and dynamical evolution of planetary systems.

SEGUE-2 will be the principal dark time project for the first year of SDSS-III. BOSS, which will use significantly upgraded versions of the SDSS spectrographs, will become the principal dark time project thereafter. MARVELS will be the primary bright time project through Spring 2011. Following completion of APOGEE's infrared spectrograph, APOGEE and MARVELS will share the bright time equally, sharing the focal plane and observing common fields where possible to maximize overall observing efficiency. In parallel with MARVELS and APOGEE, SEGUE-2 will obtain optical spectra of brighter stars ($g < 17$) in the same fields. Each of the four surveys will produce the largest and most powerful data set of its kind over the next decade, and the combination of surveys efficiently devotes all of the available observing time to a rich program of first-class science.

Like its two predecessors, SDSS-III will have an extraordinarily broad impact within the astronomical community and beyond. SDSS-III will produce large, well calibrated, easily accessible public databases that support an enormous range of astronomical research and educational activities at many levels. An active program of education and outreach will promote the SDSS-III data and tools to K-12 and university educators and to the broader public. SDSS-III will play an important role in training the next generation of outstanding young astronomers. SDSS-III focuses on some of the most exciting areas of contemporary astronomy and astrophysics, and its discoveries will have high visibility around the world.

1 Overview

The past decade has seen extraordinary progress in our understanding of the Universe, the Milky Way galaxy, and the population of extra-solar planets. A wide range of observations have converged on a consistent cosmological model that incorporates cold dark matter, a baryon-to-dark-matter ratio ∼1:6, flat space, and primordial fluctuations with the statistical properties predicted by inflation. This model requires one very surprising ingredient, "dark energy," which drives accelerating expansion of the Universe. Large-scale imaging and spectroscopic surveys have revealed rich, complex structure in the outer Milky Way, residual traces of the Galaxy's hierarchical formation history. These surveys have discovered more than a dozen new members of the Local Group of galaxies and identified thousands of the most chemically primitive stars in the Milky Way. The population of known extra-solar planets has grown from a handful to over 200. Many of the newly discovered systems are radically different from our own solar system, while others could possibly harbor life like that on Earth.

The Sloan Digital Sky Survey (SDSS), both the original 5-year program and its ongoing 3-year extension (SDSS-II), has made enormous contributions across a wide span of astronomical fields, including contributions to many of the discoveries mentioned above. Along the way it has exemplified a new mode of astronomical discovery, with teams of scientists cooperating in organized, systematic surveys to produce large data sets that are made publicly available and support a rich variety of investigations. The result has been unprecedented productivity: 1500 SDSS-based papers published in refereed journals, more than five dozen PhD theses, and 16 of the top 100 cited papers in the past five years. In a recent analysis of the observatories with the highest impact on astronomy, the SDSS has been ranked at the top for three of the last four years (see §2).

Building on this extraordinary legacy, we propose a six-year program (2008-2014) that will use the SDSS facilities at Apache Point Observatory (APO) to carry out four surveys on three scientific themes: dark energy and cosmological parameters; the structure, dynamics, and chemical evolution of the Milky Way; and the architecture of planetary systems. The wide field of view and efficient, multiplexed spectroscopic capability of the APO 2.5-m telescope make it ideally suited to exploit the new opportunities arising from recent discoveries in these fields. The overall program is divided into dark-time (moonless) and bright-time surveys. The main dark-time survey, which takes five years, is a precision measurement of the cosmic distance scale and the effects of dark energy (BOSS). The first year of dark time will be devoted to a spectroscopic study of the structure and chemical evolution of the Milky Way (SEGUE-2). The bright-time surveys are a high-resolution infrared spectroscopic survey that can penetrate obscuring dust to see deep into the heart of the Milky Way (APOGEE), and a highly multiplexed search for extra-solar planets (MARVELS). These surveys can observe simultaneously with each other and with an extension of SEGUE-2 that concentrates on bright stars. Each of the four surveys will produce the largest and most powerful data set of its kind over the next decade, and the combination of surveys efficiently devotes all of the available observing time to a rich program of first-class science.

The acceleration of the expansion of the Universe poses the most profound question in physical science today. Even the most prosaic explanations of cosmic acceleration demand the existence of a pervasive new component of the Universe with exotic physical properties. More extreme alternatives include extra spatial dimensions or a breakdown of General Relativity on cosmological scales. To distinguish competing hypotheses, we require precise measurements of the cosmic expansion history over a wide span of time. Of the four most widely studied measurement techniques, the "baryon acoustic oscillation" method, pioneered in the SDSS, is especially attractive for its simplicity and its freedom from systematic uncertainties. Sound waves that propagate in the hot plasma of the early

Universe imprint a characteristic scale on the clustering of dark matter, galaxies, and intergalactic gas. By measuring this scale with tracers seen at different redshifts, we can create a “Hubble diagram” of unprecedented precision covering most of cosmic history and can thereby pin down the properties of dark energy.

The Baryon Oscillation Spectroscopic Survey (BOSS) will map the three-dimensional distribution of 1.5 million luminous red galaxies and neutral hydrogen gas absorption in the spectra of 160,000 distant QSOs. The galaxy clustering measurements will determine the absolute distance scale with a precision of 1.0% at $z = 0.35$ and 1.1% at $z = 0.6$, extending the lever arm of the existing SDSS measurements by a factor of two and improving the measurement precision by a factor of four. BOSS will provide the definitive measurement of the low redshift ($z < 0.7$) acoustic oscillation scale, reaching close to the cosmic variance limit, and it will pioneer a powerful new method of measuring acoustic oscillations at high redshift, with the QSO absorption analysis yielding separate, 1.5%-precision measurements of the angular diameter distance and the Hubble parameter at $z = 2.5$. The BOSS measurements will achieve precise stand-alone constraints on the properties of dark energy, the curvature of space, and the Hubble constant H_0, and they will greatly strengthen the overall constraints on evolving dark energy when combined with complementary measurements that use supernovae, weak gravitational lensing, or galaxy clusters. BOSS will also provide rich insights into the matter contents of the Universe, the origin of cosmic structure, and the evolution of massive galaxies.

SDSS-I and SDSS-II have made remarkable contributions to our understanding of the Milky Way, our Galactic home. SEGUE-2 will use the SDSS spectrographs to measure line-of-sight velocities, surface temperatures, and heavy element abundances of 350,000 stars, probing the motions and chemical enrichment history of the outer Milky Way and discovering many of the lowest metallicity stars in the Galaxy. It will more than double the data set obtained in the SEGUE (Sloan Extension for Galactic Understanding and Exploration) component of SDSS-II. The APO Galactic Evolution Experiment (APOGEE) will use high-resolution ($R \sim 20,000$), high signal-to-noise ratio ($S/N \sim 100$), H-band (1.6μm) spectroscopy to penetrate the dust that obscures the inner Galaxy from our view, observing 100,000 red giant stars across the full range of the Galactic bulge, bar, disk, and halo. The high spectral resolution of APOGEE will allow element-by-element measurements of chemical abundances, which can be used to reconstruct the history of star formation that produced these elements. APOGEE will increase the number of stars observed at high spectroscopic resolution and high signal-to-noise ratio by more than a factor of 100, an extraordinary advance in the state of the art. Together, SEGUE-2 and APOGEE will provide a picture of the Milky Way that is unprecedented in scope, richness, and detail. The combined data set will play a central role in “near-field cosmology” tests of galaxy formation physics and the small scale distribution of dark matter.

Observations over the last 15 years have detected more than 200 planets orbiting around other stars. Most of these extra-solar planetary systems are very different from our own, posing major puzzles for theories of planet formation. The Multi-Object APO Radial Velocity Exoplanet Large-area Survey (MARVELS) will use fiber-fed interferometric spectrographs to monitor the radial velocities of 11,000 bright stars, with the precision and cadence needed to detect giant planets with orbital periods ranging from several hours to two years. Our forecasts predict that MARVELS will discover 150-200 new planets, mostly in the range of $0.5 - 10$ Jupiter masses. The large sample size, comprehensive coverage of stellar hosts, and well-defined statistical sensitivity will make MARVELS the critical data set for testing models of the origin and dynamical evolution of giant planet systems and the phenomenon of giant planet migration. MARVELS will complement other searches that are sensitive to low-mass or long-period planets but will not detect nearly as many giant planets in

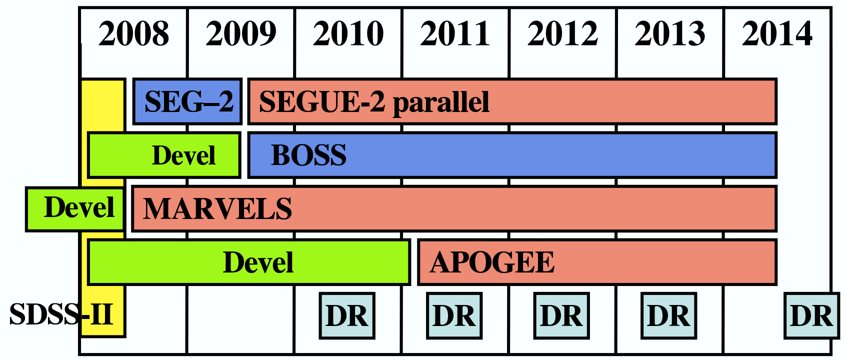

Figure 1: High-level SDSS-III schedule. Dark-time observing programs are marked in blue, bright-time observing programs (which may operate simultaneously by sharing the focal plane) are marked in red. The criterion that moonlight not significantly degrade the BOSS spectroscopic performance leaves about 60% of the observing time as "dark." Hardware development activities are marked in green, and light blue squares mark the five public SDSS-III data releases.

the dynamically evolved regime.

Figure 1 presents a high-level schedule of the four SDSS-III surveys and their development activities.

2 The SDSS Legacy

The Sloan Digital Sky Survey is one of the most ambitious and most successful projects in the history of astronomy. In its first five years of operation (2000 – 2005), the SDSS obtained five-band CCD imaging over 8,000 square degrees of the high Galactic latitude, Northern sky, detecting 217 million celestial objects. It obtained spectra of 675,000 galaxies, 90,000 quasars, and 215,000 stars, selected from 5,700 square degrees of this imaging. The data, fully calibrated and reduced, carefully checked for quality, and accessible through efficient data bases, have been publicly released in cumulative form, beginning with an early release of commissioning data and continuing with a series of annual data releases. Object catalogs, imaging data, and spectra are all available through the SDSS web site `http://www.sdss.org`, along with detailed documentation and powerful search tools.

Beginning in July 2005, the SDSS entered a new phase, SDSS-II, which will continue through June 2008. SDSS-II is using the SDSS telescope, camera, and spectrographs to carry out three distinct surveys. The Sloan Legacy Survey is completing spectroscopic observations over the full 8,000 square degrees of SDSS-I. SEGUE, the Sloan Extension for Galactic Understanding and Exploration, is imaging 3,500 square degrees on a grid extending through the plane of the Galaxy and obtaining spectra of 240,000 stars to measure their line-of-sight velocities and chemical compositions. The Sloan Supernova Survey repeatedly scans a 300 square degree area to detect and measure variable objects, with particular concentration on Type Ia supernovae to measure the expansion history of the Universe. As of this writing (August 2007), the spectroscopy of the Legacy Survey is 94% complete. SEGUE has completed 3200 square degrees of new imaging and obtained spectroscopy of 125,000 stars. In two 3-month campaigns the Supernova Survey has discovered more than 300 spectroscopically confirmed Type Ia supernovae, already exceeding its original target. Supernova candidates are announced as soon as they are discovered. The first public release from the Legacy Survey and SEGUE was SDSS DR6 in June 2007. DR6 is the first data release to incorporate "ubercalibration," which uses auxiliary observations and a novel analysis of imaging overlaps to improve the global photometric fidelity of the SDSS data by a factor of two, to approximately 1%. This high photometric precision allows the definition of accurate photometric redshifts for galaxies and the measurement of subtle variations in stellar abundances.

The SDSS data have supported fundamental work across an extraordinary range of astronomical disciplines, including the large-scale structure of the Universe, the evolution and clustering of quasars, gravitational lensing, the properties of galaxies, the members of the Local Group, the structure and stellar populations of the Milky Way, stellar astrophysics, sub-stellar objects, and small bodies in the solar system. Recent analyses by Madrid et al. (2006) and Madrid & Macchetto (2006) rate the SDSS as the most productive astronomical observatory in 2003 and 2004 based on citations to high-impact papers published in those years, ranking ahead of the European Southern Observatory (which includes the Very Large Telescope), Hubble Space Telescope, the Wilkinson Microwave Anisotropy Probe (WMAP), and the Keck Observatory. (Preliminary analyses along similar lines, privately communicated to us by R. Williams, rate the SDSS as the second most productive observatory in 2005, behind WMAP, and the most productive in 2006.) As of this writing (August 2007), searches using the NASA/ADS abstract service show that the SDSS has contributed to over 1500 refereed papers with more than 50,000 citations. Roughly half of these papers were written by astronomers outside the SDSS collaboration. Of the 100 most cited astronomical papers written since 2000 (<0.1% of the total), 16 are SDSS papers. SDSS is second only to WMAP in the past four years in most-cited papers. With the completion of SDSS-I and the broader scope of SDSS-II, the scientific impact of the SDSS seems guaranteed to increase with time.

The list of extraordinary scientific contributions of the SDSS includes, in approximately chronological order:

- *The discovery of the most distant quasars,* breaking the $z = 5$ and $z = 6$ barriers to reveal supermassive black holes in the early Universe and probe the end of the reionization epoch.

- *The discovery of large populations of field L and T dwarfs*, probing the transition from hydrogen burning stars to sub-stellar objects.

- *Mapping extended mass distributions around galaxies with weak gravitational lensing,* directly demonstrating that dark matter halos of typical field galaxies extend to 200 kpc or more and measuring the galaxy-mass correlation function for different classes of galaxies out to distances of $10 - 30$ Mpc.

- *Systematic characterization of the galaxy population,* transforming the study of galaxy properties and the correlations among them into a precise and detailed statistical science that yields powerful insights into the physical processes that govern galaxy formation.

- *The demonstration of ubiquitous substructure in the outer Milky Way,* uncovering new tidal streams and lumps and revealing striking correlations between stellar abundances and kinematics within the stellar halo.

- *Demonstration of the common origin of dynamical asteroid families* via the similar colors of objects with similar orbital properties.

- *Precision measurement of the luminosity distribution of quasars,* precisely mapping the rise and fall of quasars and the associated growth of the supermassive black holes.
- *Precision measurements of large scale clustering and cosmological constraints,* providing powerful constraints on the matter and energy contents of the Universe and on the nature and origin of the primordial fluctuations that seeded the growth of cosmic structure.
- *Precision measurement of early structure with the Lyman-α forest,* yielding precise constraints on the clustering of dark matter at redshifts of $2-4$.
- *Detailed characterization of small and intermediate scale clustering of many classes of galaxies,* allowing strong tests of galaxy formation theories and statistical determination of the relation between galaxies and dark matter halos.
- *Definition of very large samples of white dwarf stars,* including cataclysmic variables, novel objects that appear to have recently accreted asteroids or small planets, and an accurate luminosity function for testing white dwarf cooling physics and the age of the Galactic disk.
- *Discovery of many new companions of the Milky Way and Andromeda,* nearly doubling the known number of Milky Way satellites while covering just 20% of the sky.
- *Discovery of stars escaping the Galaxy,* revealing the signature of violent gravitational encounters with the Galaxy's central black hole and providing constraints on the shape, mass, and total extent of the Galactic halo.
- *Discovery of acoustic oscillation signatures in the clustering of galaxies,* opening the door to a new method of cosmological measurement.
- *Measurements of the clustering of quasars over a wide range of cosmic time,* constraining the duty cycles of active black holes and the masses of the dark matter halos that host them.

Half of these achievements were among the original "design goals" of the SDSS, but the other half were either entirely unanticipated or not expected to be nearly as exciting or powerful as they turned out to be. The SDSS and SDSS-II have enabled systematic investigation and "discovery" science in nearly equal measure, and we expect that tradition to continue with SDSS-III.

The SDSS is a valuable resource for educators at many levels, and it has made it possible for investigators with even modest computing capabilities to carry out state-of-the-art research. The SDSS has also played a major role in training new generations of outstanding scientists. More than 60 PhD theses based on SDSS data have been completed under the supervision of SDSS participants (and probably many more with public data). Equally striking, six of the ten members of the current SDSS-III Management Committee, including the Director, the Project Scientist, and the BOSS and SEGUE-2 PIs, originally joined the SDSS project as postdocs (or, in one case as an undergraduate), rising through the ranks to leadership roles.

When SDSS-II is completed in mid-2008, the SDSS facilities will remain a uniquely powerful resource for wide-field spectroscopic surveys. (The SDSS camera, while still among the world's most powerful, is no longer unique, which is why SDSS-III emphasizes spectroscopic programs.) While the telescope and instruments are extraordinary assets in themselves, the technical infrastructure and collaboration culture that transform raw observations into calibrated, accessible data and high-impact science are equally important. SDSS-III will extend the legacy and leverage the accomplishments of the Sloan Digital Sky Survey.

3 BOSS: Probing Dark Energy and Cosmological Physics

The past decade has been one of extraordinary progress in cosmology, with perhaps the most startling discovery being that the expansion of the Universe is accelerating. Such acceleration poses a deep challenge to theories that seek to understand the fundamental nature of space and time, and the basic constituents of matter. Theoretical ideas for explaining cosmic acceleration, which we shall refer to collectively as dark energy, include Einstein's cosmological constant (Einstein, 1917), which can be interpreted as zero-point energy of the quantum vacuum, a new scalar field (Ratra & Peebles, 1988; Wetterich, 1988; Coble et al., 1997) with either canonical or non-canonical (Armendariz-Picon et al., 2000) kinetic terms, holographic dark energy (Bousso, 2002), modifications of General Relativity such as scalar-tensor theories or $f(R)$ theories (Capozziello, 2002; Carroll et al., 2004), or the signature of extra dimensions as in braneworld cosmologies (Arkani-Hamed, Dimopoulos & Dvali, 1998). No current theory gives a compelling explanation of why the dark energy density is the observed magnitude or why it happens to be close to the matter density today.

Going beyond the detection of acceleration to informative constraints on its origin requires measurements of cosmic expansion with percent-level precision and exquisite control of systematic uncertainties. At present, there are four main observational probes of dark energy (Albrecht et al., 2006), which aim to measure distances as a function of redshift, the growth of structure, and possible fluctuations in dark energy at large scales. Type Ia supernovae, which provided the first direct evidence for cosmic acceleration (Riess et al., 1998; Perlmutter et al., 1999), measure luminosity distance vs. redshift and thus provide a purely geometrical constraint. The evolution of the abundance of rich clusters depends on both the geometry and the growth of perturbations, as does weak gravitational lensing. A fourth method, also purely geometrical, uses a standard ruler to measure the angular diameter distance vs. redshift. It is this fourth method, using baryon acoustic oscillations (described further in §3.1), which we aim to exploit.

At present, baryon acoustic oscillations (BAO) are believed to be the method "least affected by systematic uncertainties, and for which we have the most reliable forecasts of resources required to accomplish a survey of chosen accuracy" (report of the Dark Energy Task Force; Albrecht et al. 2006). Sound waves that propagate in the opaque early universe imprint a characteristic scale in the clustering of matter, providing a "standard ruler" whose length can be computed using straightforward physics and parameters that are tightly constrained by cosmic microwave background (CMB) observations. Measuring the angle subtended by this scale determines a distance to that redshift and constrains the expansion rate. The detection of the acoustic oscillation scale is one of the signature accomplishments of the SDSS (see Fig. 2), and even this moderate signal-to-noise measurement substantially tightens constraints on cosmological parameters (Eisenstein et al., 2005; Hütsi, 2006; Tegmark et al., 2006; Padmanabhan et al., 2007; Percival et al., 2007a,b).

The challenge for the BAO technique, also noted by the DETF, is that very large surveys are required to attain high statistical precision. The Baryon Oscillation Spectroscopic Survey (BOSS) will survey the immense volume required to obtain percent-level measurements of the BAO scale. BOSS is a next-generation experiment that will transform the BAO technique into a precision cosmological probe by leveraging the proven hardware and the experience of a collaboration that made one of the first measurements of BAO. A natural extension of the SDSS and SDSS-II redshift surveys, the BOSS project has two distinct components that are carried out simultaneously: a galaxy redshift survey of 1.5 million luminous red galaxies (LRGs) at $0.2 < z < 0.8$ over one quarter of the sky and a survey of 160,000 quasi-stellar objects (QSOs) at $2.3 < z < 2.8$ that will allow us to probe the intergalactic medium along each QSO sightline. Our forecasts show that the LRG and QSO surveys will yield absolute distance measurements with precision of 1.0% at

$z = 0.35$, 1.1% at $z = 0.6$, and 1.5% at $z = 2.5$, greatly tightening the constraints on dark energy and the curvature of space (Fig. 5 and Table 1; discussed further in §3.4).

The BOSS LRG survey will be the definitive low redshift ($z < 0.7$) BAO experiment for the foreseeable future because it covers a large area of sky with precise spectroscopic redshifts of strongly clustered tracers, with sampling density sufficient to limit shot-noise contributions to statistical errors. In contrast to Type Ia supernovae, which are calibrated in the local Hubble flow, the leverage of BAO measurements on dark energy parameters is strongest at *low* redshift, because the absolute distance scale is anchored in the cosmic microwave background. BOSS will also lay the groundwork and provide the essential low redshift comparison point for future baryon oscillation experiments at higher redshifts. The BOSS QSO survey will itself pioneer a novel method for high-redshift BAO measurements, achieving valuable constraints on its own and providing a pathfinder for still more powerful surveys that could use this technique in the future. BOSS will complement other future experiments that use different observational probes, such as supernovae, clusters or weak lensing (see §3.4). Finally, the high precision, enormous dynamic range, and wide redshift span of the BOSS clustering measurements will greatly improve empirical constraints on neutrino masses, the physics of inflation, the evolution of galaxies, and the physics of supermassive black-holes powering QSOs.

All of this can be accomplished with readily achievable improvements to the SDSS spectrographs and fiber system and a slight increase of integration times relative to the SDSS. During Fall 2008, BOSS will use the SDSS camera to carry out an additional 2,000 square degrees of imaging at high Galactic latitude in the southern Galactic cap, filling the gaps between the SDSS southern stripes. After the spectrograph upgrades are completed in Summer 2009, BOSS will carry out a redshift survey of 1.5 million LRGs over $10,000$ square degrees and $160,000$ QSOs over a somewhat smaller $8,000$ square degree footprint. The latter will be observed only in dark time, while the former will be observed in both dark and grey time.

3.1 Baryon acoustic oscillations

Baryon acoustic oscillations arise due to the tight coupling of baryons and photons in the hot, dense early Universe. During tight coupling, perturbations in the baryon-photon fluid are unable to grow; instead, they propagate as sound waves with a speed slightly over half the speed of light (Peebles & Yu, 1970; Sunyaev & Zel'dovich, 1970). Soon after recombination ($z \simeq 10^3$, $t \simeq 400,000\,\mathrm{yr}$) the baryons and photons decouple, allowing the perturbations to begin growing by gravitational instability. The sound waves are frozen into the plasma at decoupling, and their imprint can be detected in the late-time distribution of matter as an enhancement of clustering at 150 Mpc separation or a series of "wiggles" in the power spectrum (Peebles & Yu 1970; Sunyaev & Zel'dovich 1970; see Eisenstein, Seo & White 2007 for a pedagogical description). The subsequent evolution of the matter slightly smears this feature (or suppresses the higher harmonics in Fourier space) but leaves it otherwise intact (see Figure 3). Because the length of the acoustic scale is precisely constrained by measurements of CMB anisotropy, we can use the galaxy measurements of the acoustic peak as a standard ruler (Eisenstein, Hu & Tegmark, 1998; Eisenstein, 2003; Blake & Glazebrook, 2003; Hu & Haiman, 2003; Seo & Eisenstein, 2003). Indeed, by measuring the acoustic scale along and across the line of sight one can measure two cosmological distances, the Hubble parameter $H(z)$ and the angular diameter distance $d_A(z)$, at the redshift(s) of the survey. As they are tied to the CMB, the calibration of these measurements in absolute units (e.g., meters or Mpc) is independent of the Hubble constant, in contrast to measurements based on supernovae or other distance indicators calibrated in the local Hubble flow. The SDSS has achieved a 4% measurement

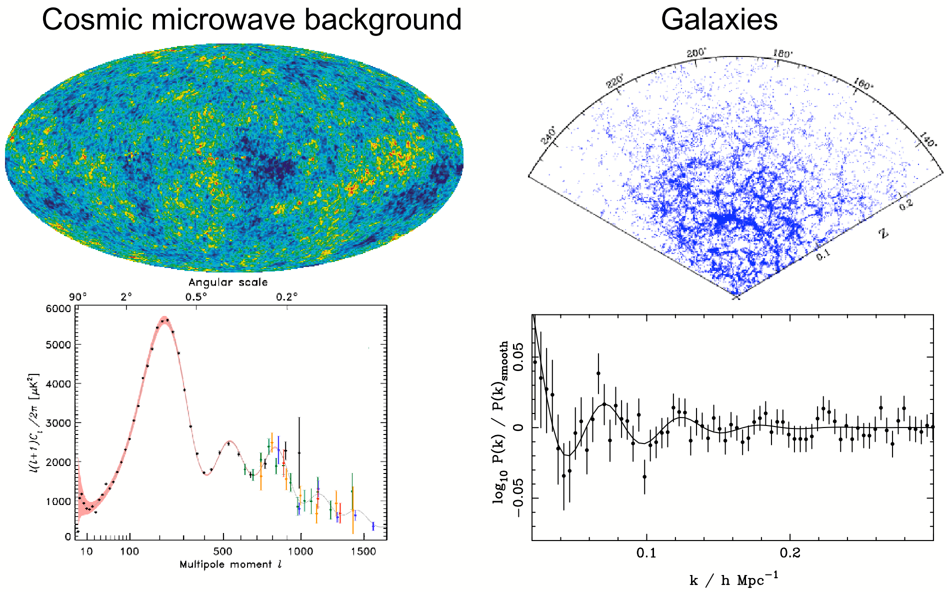

Figure 2: Acoustic oscillations seen in the very early Universe and in the present day distribution of galaxies. The top left panel is a map of the fluctuations in the CMB (from Kogut et al., 2007), which show the imprint of acoustic waves – the series of peaks in the lower left panel – when the Universe was $400,000$ years old. The top right panel shows the distribution of galaxies in the SDSS survey, and the lower right panel shows that these have the same scale imprinted upon them (from Percival et al., 2007a).

of the distance to $z = 0.35$ (Eisenstein et al., 2005; Tegmark et al., 2006; Percival et al., 2007a) and 6% to $z = 0.5$ (Blake et al., 2007; Padmanabhan et al., 2007).

Because the acoustic scale is so much larger than the scale of non-linear gravitational collapse, the method is highly robust. There is an increasing body of analytic and numerical work (Meiksin et al., 1999; Seo & Eisenstein, 2005; Springel et al., 2005; White, 2005; Eisenstein, Seo & White, 2007; Huff et al., 2007; Angulo et al., 2007; Nishimichi et al., 2007) arguing that the derived distances are unbiased at the 1% level, once the smearing of the peak is accounted for and broad-band tilts are removed. Eisenstein, Seo & White (2007) argue that the shift should be 0.5% for reasonable galaxy populations. Crocce & Scoccimarro (2007) argue that second-order terms in the matter also yield shifts at or below 0.5%; this shift has yet to be detected in any simulation, but work is ongoing. Even if continuing work shows that non-linear physics causes percent-level shifts in the acoustic scale, we need only calculate these non-linear corrections with moderate accuracy to bring systematic errors well below the level of our statistical uncertainties. Careful peak centroiding algorithms can remove broad-band trends caused by non-linear evolution or scale-dependent galaxy bias, which could otherwise bias the peak position (Smith, Scoccimarro & Sheth, 2007). Such detrending algorithms can also take care of higher order effects, such as gravitational lensing (Vallinotto et al., 2007; Hui, Gaztanaga & LoVerde, 2007). The 'nuisance' parameters associated with such methods couple very weakly to the cosmological information of interest, so marginalizing over them has a negligible effect on our distance errors (see references above).

The smearing of the peak in the galaxy correlation function $\xi(r)$ is dominated by the formation

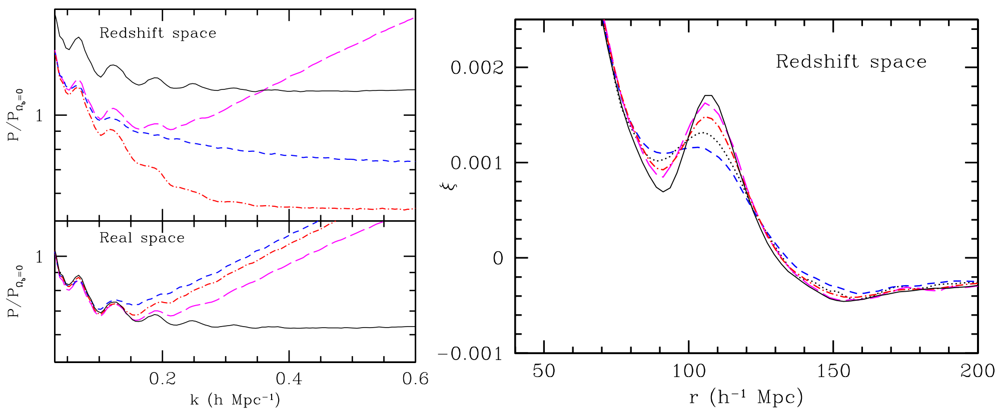

Figure 3: (Left) The linear theory matter power spectrum (black, solid; displaced upwards in top panel for clarity), the spectrum at $z = 0.3$ (blue, short-dashed) which has had the peaks partially erased by non-linear evolution, and the spectrum recovered from reconstruction (red, dot-dashed and magenta long-dashed) with two different filters. Note that the peaks at low k are restored – the difference in power at high k would be marginalized over in our fits. (Right) The same effects in the correlation function, where the broadening of the peak by non-linearity and the gains from reconstruction are evident. Our signal comes from centroiding this peak. From Eisenstein et al. (2007a).

of superclusters and voids — large scale gravitational sources. Given a good enough measurement of the density field one can undo some of the non-linear 'smearing' of the acoustic feature (Eisenstein et al., 2007a). By moving galaxies back along the velocity (or displacement) vectors inferred from their density field, one sharpens the acoustic peak (Fig. 3), allowing more precise distance measurements. In essence, non-linear evolution moves information out of the two-point function into higher order correlations, or equivalently into Fourier phase correlations. Even a simple massaging of the density field can bring some of that information back into the power spectrum. The reconstruction need not be perfect – fixing only half of the non-linear displacement substantially improves recovery of the acoustic scale at low z. Since such a reconstruction only relies on scales larger than $20\,h^{-1}$Mpc, the task is much easier than 1990s-era work on velocity fields, which used much smaller surveys and focused on smaller scales.

In principle, BAO measurements can be done with photometric redshifts, but with significant loss in constraining power. A photo-z survey over the same area as BOSS to similar depth would yield roughly 3% constraints on d_A, instead of BOSS's 1%, because the acoustic oscillation peak would not be resolved in the redshift direction (e.g., Seo & Eisenstein, 2007). Even a full-sky photo-z survey to the same redshift would not provide the statistical precision of BOSS. Furthermore, a spectroscopic map greatly reduces calibration-related systematic uncertainties in the distance scale, offers the ability to constrain *both* $H(z)$ and $d_A(z)$, and opens the possibility of reconstruction to mitigate the effects of non-linearity at lower z. The wide-field spectroscopic capabilities of the APO 2.5 m telescope are ideally suited to this task.

BOSS will use the two probes of high-redshift structure that were proven so successful in SDSS-I: the clustering of LRGs and structure in the Ly-α forest of $z > 2.3$ QSOs. As in SDSS-I, we will use the multi-color SDSS imaging to select the high redshift LRGs. Our group already has experience selecting distant LRGs from the 2dF-SDSS LRG and Quasar survey (2SLAQ; Cannon et al. 2006) sample, which observed over 6,000 LRGs at similar redshifts to the BOSS sample. Using existing

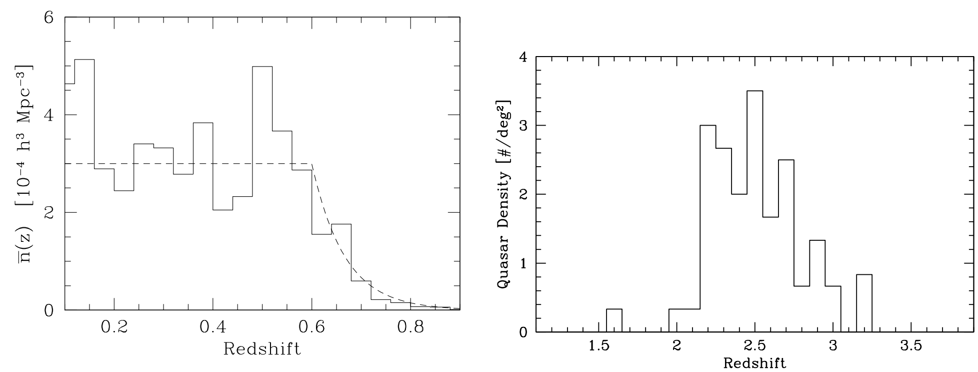

Figure 4: The redshift distribution of LRGs (left) and QSOs (right) for our baseline target selection. These targets are selected from SDSS imaging over 7.6 deg^2 with spectroscopic redshifts from AGES (`http://cmb.as.arizona.edu/∼eisenste/AGES/`). Due to the small field one can see large-scale structure in the histograms; nevertheless the figure demonstrates that we can achieve our target number density and redshift distribution from SDSS imaging.

data over a small patch of sky, we have demonstrated that BOSS can achieve the desired redshift distribution and areal density using techniques similar to the original SDSS LRG selection, but now reaching a depth of $i \simeq 20$ (compared to $r = 19.5$, hence $i \simeq 18.5$, in SDSS-I/II). With improved spectrograph performance the desired number density can be reached out to $z = 0.5$, with useful galaxies out to $z = 0.8$. The baseline redshift histogram is shown in Fig. 4. By covering a larger volume and sampling more densely, we improve the SDSS-I measurement at low z (the SDSS-I LRG results had a significant shot-noise contribution to the error bar) and extend it to higher z (Fig. 6). The higher sampling density also enables the kind of reconstruction illustrated in Fig. 3, as shown by Eisenstein et al. (2007a).

We will also use the Ly-α forest in bright QSO spectra to probe the density fluctuations at high redshift (White, 2003; McDonald & Eisenstein, 2006). The Ly-α optical depth is closely tied to the underlying dark matter density (Cen et al., 1994; Croft et al., 1998; McDonald et al., 2000), so correlations in the flux can be used to measure correlations in the dark matter. Using 3000 SDSS QSO spectra probing the redshift range $2.2 \leq z \leq 4$, McDonald et al. (2005) obtained precise constraints on the shape and amplitude of the high-redshift matter power spectrum, which Seljak et al. (2005) used to obtain tight cosmological parameter constraints. Each QSO spectrum gives data on hundreds of points in space, allowing BOSS to measure the angular diameter distance to and Hubble parameter at $z \sim 2.3$. The dense grid of 160,000 QSO sightlines in BOSS will enable measurements that are effectively three-dimensional, instead of treating each sightline independently. The gain in statistical power is thus much larger than the "mere" 50-fold increase in sample size would suggest, and, critically, the measurement becomes insensitive to continuum determination errors, which should be uncorrelated across sightlines. The required spectral resolution is well below that of the SDSS spectrographs, and one can use spectra of even modest signal-to-noise ratio. This allows the QSO survey to be carried out in parallel with the LRG survey using the same spectroscopic plates and exposures.

3.2 Instrument upgrade

The SDSS telescope still has enormous benefits over any other spectroscopic facility because of its 7 $\deg^2$ field of view. It has successfully measured redshifts for $\sim$ 99% of the targeted LRGs out to $z = 0.45$. The instrumental capabilities that we plan to improve are (a) the number of fibers per field, (b) the red sensitivity and UV throughput of the CCDs, and (c) the size of the fibers. One thousand fibers with 2 arcsec diameters will replace the current 640 fibers with 3 arcsec diameters and will focus to larger CCDs with smaller pixels. All BOSS target fibers will be devoted to LRGs or high-redshift QSOs, while in SDSS-I and SDSS-II only 15% of galaxy targets were LRGs at $z > 0.15$ and only 12% of QSO targets were at $z > 2.3$ as required for Ly-α forest measurement. The spectrographs need to be sufficiently red-sensitive to measure the 4000 Å break and the strongest Mg lines redshifted to $\sim$ 9300Å. The improved wavelength coverage will be 3700 Å - 9800 Å, with a resolution on the red side increasing modestly from the current $R = 1800$ with 2.4 pixels per resolution element to $R = 2400$ with 3 pixels per resolution element.

Our design maintains those parts of the spectrographs that are sufficiently efficient and have proven so successful. We will retain the fiber plugging as it is done now but replace the fibers in all cartridges. We will switch to thick, fully depleted CCDs which have improved quantum efficiency in the near-infrared and much reduced fringing (Bebek et al., 2004). For this project, LBNL will provide two 4K$\times$4K devices similar to those used at Lick observatory and proposed for the Dark Energy Survey. LBNL has fabricated one lot (24 wafers, 21 potential CCDs) of fully-depleted red 4K$\times$4K CCDs at DALSA, and they are now at LBNL for backside processing, packaging and testing. The blue-side CCDs will be swapped with off-the-shelf Fairchild 4K x 4K devices with 15μm pixels.

We will improve the optics by replacing an optical coupling grease between two lenses in the spectrograph cameras, and by replacing the gratings with volume phase holographic (VPH) grisms. The peak throughput will be increased by 50%, with more improvement at the scientifically critical wavelengths of the UV and near-IR.

3.3 BOSS in context

The Dark Energy Task Force (DETF; Albrecht et al., 2006) made specific recommendations for the next round of dark energy experiments, which they dubbed Stage III, following on the discovery experiments of Stage I and the Stage II experiments that are currently in progress. For Stage III, the DETF recommended that funding agencies select a coordinated set of experiments that include all four techniques: Type Ia supernovae, weak lensing, clusters, and baryon oscillations. BOSS is a prototypical Stage III experiment and occupies a unique place in the cadre of those that are likely to go ahead.[2] In particular, BOSS is an ideal complement to two other major Stage III efforts, the Dark Energy Survey (DES; Abbott et al., 2005) and Pan-STARRS (`http://pan-starrs.ifa.hawaii.edu/`), which focus on shear-shear weak lensing, supernovae, and clusters.

BOSS uses spectroscopic redshifts to achieve the full power of BAO; photometric redshifts provide a washed-out BAO map that is less statistically constraining, more subject to systematic uncertainties, and yields no direct measurement of $H(z)$. With its 10,000 square degree survey

[2]The call for DETF White Papers on experimental approaches to dark energy closed in June 2005, before SDSS-II had even begun, and it was not until Fall 2005 that we began to think in earnest about possible post-2008 programs with the 2.5-meter telescope. Therefore, BOSS does not appear in the DETF report. However, it yields more powerful dark energy constraints than the Stage III BAO experiments considered by the DETF (see Table 1 below).

area, BOSS is within a factor of two of the cosmic variance limit for BAO measurement precision at $z < 0.6$, so no future experiment can do substantially better in this regime. It is the most powerful BAO experiment planned for the next decade (see below).

BAO is a powerful technique on its own, but it is even more powerful in combination with Type Ia supernovae, the most established method and the one that currently contributes the most to dark energy constraints (Albrecht et al., 2006). By tying SNe distances to the physical scale provided by the CMB, and by extending the distance lever arm to high redshift, the BAO technique naturally complements SNe measurements. As shown in the next section, BOSS in combination with the DETF Stage III supernovae and existing Stage II constraints yields an improvement by a factor of more than three in the DETF figure-of-merit (defined below) over Stage II alone. Because BAO measurements are purely geometrical, they complement weak lensing and cluster measurements that depend on both geometry and the growth of structure.

BOSS is cost-effective and timely, building on the substantial infrastructure of SDSS-I and II. Timeliness is important because the DETF recommended that significant weight be placed on the ability of Stage III experiments to inform the design of future Stage IV experiments. BOSS will test the potential of Stage IV BAO experiments such as the ground-based Square Kilometer Array (SKA) or space-based ADEPT satellite, and it will provide a low-redshift anchor for ADEPT if it goes ahead. In addition, the QSO component of BOSS could demonstrate a new way to measure BAO at high z, significantly impacting the design of future BAO experiments in this regime. BOSS will also cross-check the SNe distances from other Stage III experiments, validating them at high accuracy before Stage IV experiments begin. BOSS would be fully complementary to the Large Synoptic Survey Telescope (LSST), which would obtain precise BAO measurements at intermediate redshifts using photometric redshifts (in addition to its supernova, weak lensing, and cluster measurements) but would not supersede BOSS at $z < 0.7$ or at $z \sim 2.5$.

3.4 BOSS science

The observed cosmic acceleration has presented us with many puzzles. Is it caused by a breakdown of General Relativity on cosmological scales or by a mysterious dark energy pervading all space? If the latter, is the dark energy a vacuum energy (cosmological constant) or an evolving field? If the latter, what is its equation-of-state ($w \equiv P/\rho$), and has this evolved with time? BOSS will directly test the cosmological constant and evolving scalar field hypotheses, and it will provide crucial constraints that can be combined with other measurements to test General Relativity. Along the way it will necessarily address some closely intertwined questions: is space curved, what is the expansion rate today (the Hubble constant), and what is the matter density (Ω_m)?

To characterize the capabilities of BOSS, we follow common practice and parameterize $w(a) = w_0 + w_a(1-a) = w_p + w_a(a_p - a)$, where $a \equiv (1+z)^{-1}$ and a_p is a 'pivot' point where $w(a)$ is best constrained and the errors on w_p and w_a are uncorrelated. The Hubble parameter $H(a)$ is then

$$H^2(a) = \Omega_m a^{-3} + \Omega_r a^{-4} + \Omega_K a^{-2} + \Omega_X a^{-3(1+w_0+w_a)} e^{-3w_a(1-a)} \quad , \tag{1}$$

where Ω_i is the present fraction of the critical density in species i and X denotes dark energy. The angular diameter distance at redshift $z = a^{-1} - 1$ is then

$$d_A = |K|^{-1/2} \mathrm{sinn}\left(|K|^{1/2}\chi\right) \quad , \qquad \chi = \int_a^1 \frac{da'}{a'^2 H} \quad , \tag{2}$$

where $\mathrm{sinn}(x) = \sin x$, x or $\sinh x$ for closed, flat or open cosmologies and $|K|$ is the spatial curvature.

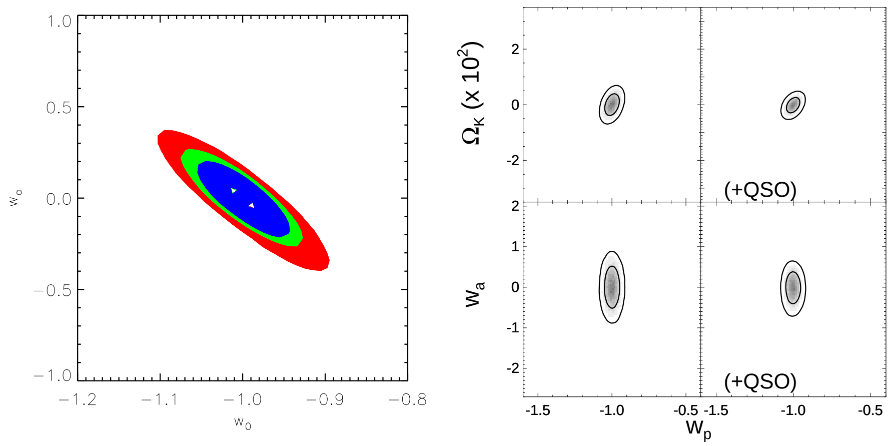

Figure 5: (Left) Red and blue contours show the 68% marginalized constraints on the DE parameters w_0 and w_a from combining BOSS BAO constraints with the DETF constraint forecasts for, respectively, Stage II experiments and Stage II + weak lensing, cluster and supernovae Stage III experiments. Green contours show the DETF forecasts for Stage II + Stage III without BOSS. (Right) The contours for Stage II + BOSS BAO in the Ω_K, w_p, w_a plane, with and without the QSO component.

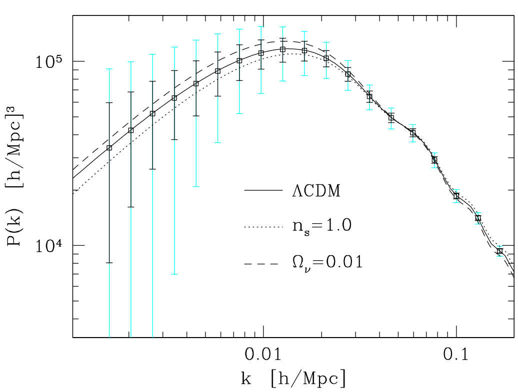

Figure 6: Forecast errors on the LRG power spectrum from BOSS (points with dark error bars) and a final analysis of $8,000\,\mathrm{deg}^2$ from SDSS-I (light error bars). Also shown are theoretical predictions for models with scale-invariant spectra (instead of $n_s = 0.97$) or massive neutrinos (instead of $\Omega_\nu = 0$); parameter changes of this magnitude could be ruled out, as the statistical errors are only weakly correlated.

In Fig. 5 and Table 1 we forecast the constraints on a number of cosmological parameters. To obtain these numbers we first convert our observational parameters into errors on the line-of-sight $[H(z)]$ and transverse $[d_A(z)]$ distances as a function of redshift using the method of Seo & Eisenstein (2007). This Fisher matrix calculation uses only acoustic oscillation information and no broad-band power, so we believe the error estimates to be robust (and conservative, see the discussion below and the final two lines of Table 1). To approximate the effects of (partial) reconstruction (Eisenstein et al., 2007a) we suppress the nonlinear smearing ($\Sigma_{\perp,||}$ in the notation of Eisenstein, Seo & White 2007) by a factor of two for the LRG calculation. We use a similar Fisher matrix calculation (McDonald & Eisenstein, 2006) to estimate the distance errors that one would obtain for the QSOs, with no attempt at reconstruction because of the very low number density. We find errors on d_A of 1.0% at $z = 0.35$, 1.1% at $z = 0.6$, and 1.5% at $z = 2.5$, with errors on $H(z)$ of 1.8%, 1.7% and 1.5% at the same redshifts. As noted earlier, current theoretical studies suggest that any shifts in the BAO scale due to non-linearity or galaxy bias are at or below this level. With further work, we should be able to calculate any corrections to a level of accuracy that keeps systematic errors well below these statistical errors.

The constraints on d_A and H are then used in a Markov chain to get constraints on the matter density $\omega_m \equiv \Omega_m h^2$, the baryon density $\omega_b \equiv \Omega_b h^2$, the dark energy density Ω_X, w_0, w_a, and the curvature, Ω_K. In addition to the distance constraints from BAO experiments, we apply a (Gaussian) prior on these parameters equal to the Fisher matrices for Planck, HST, and Stage II experiments presented in the technical appendix of the DETF report. The Markov chain allows us to explore the full likelihood surface, without making a "quadratic" assumption like Fisher matrix methods do. This is important when some parameters (e.g. w_a) are not well constrained. The variance of each parameter, computed from the chains, is given in Table 1. We also quote the DETF Figure of Merit, which is the inverse of the area of the 95% confidence level region in the $w_p - w_a$ plane (scaled to correspond to the convention adopted by the DETF). The precise value of the pivot expansion factor depends on which experiments are considered, but it is generally $a_p \approx 0.8$, i.e., in this family of models the dark energy experiments best constrain the value of w at $z \approx 0.25$.

Table 1 makes clear that BOSS will provide an impressive advance over our current and near-future state of knowledge. The QSO survey substantially tightens the parameter constraints, in large part because it sharpens the constraint on space curvature. The comparison with other proposed BAO experiments is given in the Table; BOSS is the most powerful experiment as measured by the FoM, with only WFMOS (which would come much later than BOSS) coming close. These comparisons are difficult because not all of the experiments have frozen their design. For WiggleZ we use the description in Glazebrook et al. (2007), for HETDEX we use Hill et al. (2004), and for WFMOS Eisenstein et al. (2007b). It is possible that some of the experiments' plans could change with further optimization. It is even harder to compare the BOSS performance to that of other methods because of the difficulty in accounting for method- or experiment-specific systematics. However, we note that the FoM forecast for BOSS is similar to the individual DETF forecasts for other Stage III methods, which assume survey parameters comparable to those planned for DES and Pan-STARRS. Since the BOSS measurement is independent of the cluster, weak lensing, and supernova measurements, it allows a cross-check of dark energy conclusions at a similar signal-to-noise ratio. The combination of two measurements with similar signal-to-noise ratio more than doubles the FoM because the degeneracy directions of the constraints are different. For example, the strong BOSS constraint on spatial curvature breaks a key degeneracy in dark energy parameter constraints for SNe experiments, even at high z. The combination of BOSS and Stage II constraints increases the FoM by a factor of 2.3 over Stage II alone, while the combination of BOSS with Stage

Expt.	h	Ω_K	w_0	w_p	w_a	FoM
BOSS LRG	0.008	0.0028	0.089	0.032	0.366	86
BOSS LRG+QSO	0.008	0.0019	0.076	0.029	0.279	**122**
+WL	0.008	0.0017	0.068	0.026	0.227	172
+CL	0.008	0.0018	0.071	0.023	0.244	177
+SN	0.006	0.0019	0.052	0.023	0.220	199
+WL+CL+SN	0.005	0.0016	0.046	0.018	0.164	331
WiggleZ	0.012	0.0028	0.099	0.035	0.430	66
HETDEX	0.015	0.0021	0.098	0.034	0.417	70
WFMOS	0.011	0.0017	0.083	0.033	0.323	95
Including Broad-Band Power Information:						
BOSS LRG+QSO	0.007	0.0015	0.065	0.016	0.240	257
+WL+CL+SN	0.005	0.0013	0.041	0.014	0.150	479

Table 1: A comparison of the abilities of current and next generation BAO experiments to constrain the expansion rate and curvature of space and the redshift dependent equation of state of dark energy. All constraints assume the DETF forecasts for "Stage II" experiments, which alone have a FoM of 53, as a prior. Lines $3-6$ show the additional gains from adding Stage III weak lensing, cluster, and supernova constraints, using the DETF "optimistic" forecasts. BAO constraints in the first two sections of the table include only the acoustic scale information and are therefore conservative; the final two lines show BOSS forecasts that also incorporate broad-band power information. See text for details.

II and the Stage III DETF forecasts for other methods yields a factor of 6.2 increase. (We have used the "optimistic" DETF Fisher matrices for other experiments.)

The BOSS forecasts in Figure 5 and the top section of Table 1 are conservative because they include only the BAO measurements, not the additional information in the shape and anisotropy of the broad band power spectrum. The anisotropy, which can be measured at high signal-to-noise over a wide range of scales, provides constraints on spacetime geometry (Alcock & Pacynski, 1979) and, via peculiar velocity distortions, on $\Omega_m(z)$ (Kaiser 1987). We have also computed Fisher matrix forecasts based on linear theory that incorporate the full LRG power spectrum information, cutting off at the non-linear scales calculated by Eisenstein, Seo, & White (2007); these are reported in the last two lines of Table 1. The error bar on w_p improves by a factor of 1.8, and the FoM grows by a factor of two. Extracting these more aggressive constraints in practice will require advancing the state of the art in theoretical modeling of non-linear gravitational evolution and galaxy bias, so it is presently difficult to project the associated systematic uncertainties. However, we believe that these challenges are no greater than those required to achieve the "optimistic" (or even middle-of-the-road) level of systematic errors in Stage III experiments based on weak lensing or galaxy clusters. Furthermore, we have not considered the additional gains (or at least cross-checks) that can be obtained using the bispectrum or other higher order clustering statistics.

While the FoM performance of BOSS is impressive, this measure alone understates the true value of BOSS because it presumes that dark energy is accurately described by the (w_0, w_a) parameterization. Clear evidence for *departures* from this model would provide direct insight into the origin of cosmic acceleration. The discovery of such departures, if they exist, will rely crucially on combining measures with different information content. For example, a comparison of BOSS distance-redshift results with the cluster and weak lensing results would be a powerful test of General Relativity, since the latter depend on the gravitational growth of inhomogeneities while the former does not. Specifically, if we assume that General Relativity and the (w_0, w_a) parameterization are correct, we find that the BOSS geometrical constraints allow us to predict the ratios $D(z=2)/D(z=1000)$ and $D(z=0.5)/D(z=1000)$ to 0.6% and 1.0%, respectively, where $D(z)$ is the linear growth factor at redshift z. These ratios can be measured directly by combining CMB data with cluster or weak lensing data, allowing the consistency of General Relativity on cosmological scales to be tested at the 1% level.

The direct measurement of $H(z)$ at $z \sim 2.5$ from the QSO survey will allow a key test of "early dark energy models," in which the relative contribution of dark energy at high redshifts is larger than it is in (w_0, w_a) models (e.g., Albrecht & Skordis, 2000). This class of models offers one route to ameliorating the "coincidence problem," the similarity of present day matter and dark energy densities. Because $H^2(z)$ is proportional to the energy density, the BOSS QSO survey could detect few percent departures from the energy scaling predicted by standard models. Substantial evidence for such departures would transform the approach to characterizing dark energy and guide the design of Stage IV BAO experiments that could measure them at higher precision.

Another remarkable outcome of BOSS will be a determination of the Hubble constant H_0 with a precision of better than 1%, a factor of ~ 10 improvement over the HST Key Project determination (see Table 1). This determination becomes possible because the BAO measurements yield an absolute calibration of the supernova distance scale, in Mpc, which is then transferred to low redshift by the Stage II supernova experiments. Combination of BOSS with the SDSS-II supernova survey, which probes the redshift range $z = 0.1 - 0.4$, would allow this determination to be based on just two experiments. The combination of BAO and supernovae will invert the traditional cosmic distance ladder, with the absolute luminosities of stellar "distance" indicators inferred at the 1% level from the known value of H_0.

The strong BOSS constraint on spatial curvature will provide a valuable test of inflation models, which generically predict that Ω_K is immeasurably different from zero. A measurement establishing $|\Omega_K| = 0.01$, say, would require either fine-tuned inflationary models or an alternative explanation for the near but not perfect flatness of the Universe. If the Universe is homogeneous on superhorizon scales and is simply connected topologically, then the sign of Ω_K determines whether space is finite or infinite; a measurement of non-zero Ω_K would inform theoretical ideas about quantum cosmology and the nature of the Big Bang itself. Note that the tight constraints on Ω_K quoted for many experiments, particularly CMB experiments, generally *assume* that dark energy is a cosmological constant, and they weaken for more general dark energy models. The combination of CMB data with a direct measurement of $d_A(z=2.5)$ from the BOSS QSO survey yields a tight constraint on Ω_K with little sensitivity to dark energy uncertainties.

An additional route to constraining dark energy and structure formation comes from a galaxy-galaxy lensing analysis (Hirata & Seljak, 2003; Sheldon et al., 2004) of the BOSS LRG sample. This will allow a determination of the correlation between LRGs and dark matter over a wide range of scales. Accounting for the sample size and redshift distribution of BOSS we expect to achieve a $S/N = 200$ measurement of the LRG-mass cross-correlation. The small-scale signal will allow us to measure dark matter halo profiles to unprecedented precision while the large-scale signal, under the assumption that LRGs trace the dark matter with some bias, allows us to constrain the dark matter auto-correlation function at the LRG redshift. This method is complementary to the weak lensing method of shear-shear correlations pursued by other Stage III projects. The availability of spectroscopic redshifts for lenses combined with well calibrated photometric redshifts for sources allows for some unique advantages over these deeper shear-shear analyses: it typically has dramatically higher S/N, it allows one to measure the dark matter clustering amplitude as a function of z rather than just in broad windows in z, and it does not suffer from some of the intrinsic alignment effects. We anticipate a 1-2% statistical error on the amplitude of matter fluctuations. The method can be easily extended once deeper imaging surveys such as Pan-STARRS or LSST become available over the same area, highlighting the synergy between these different surveys.

BOSS would be the largest effective volume yet surveyed for large-scale structure. Including the effects of shot noise, the LRG survey will measure a quarter million Fourier modes at $k < 0.2\,h\mathrm{Mpc}^{-1}$, over 7× the number in SDSS-II (Fig. 6). The BOSS Ly-α forest spectra would significantly increase the total line-of-sight distance available for study, providing strong constraints on small-scale power. While the most important result from the power spectrum will likely be the acoustic scale discussed above, the BOSS spectrum will also be a powerful cross-check on the Planck cosmology. CMB and large-scale structure data together can be used to test the primordial power spectrum in a model-independent fashion, allowing tight constraints on inflation models and neutrino masses and high-precision tests of our standard account of the growth of cosmic structure. For example, using standard Fisher matrix methods we forecast that a combination of Planck CMB data and BOSS can halve the present upper bound on the (summed) neutrino mass, even with conservative assumptions about modeling galaxy bias. Any discrepancies between Planck CMB data and BOSS could point to important physics missing in the standard cosmological model.

The BOSS data set will be an enormous sample in which to study the evolution of massive galaxies, traditionally a challenge for theoretical models, which tend to make these galaxies too massive and too blue. The combination of SDSS and BOSS data, along with other surveys, will allow us to trace the process of galaxy formation over half the age of the Universe. Since activity levels in massive systems are expected to increase at earlier epochs, such a lever arm in time promises to teach us how the galaxies have been assembled, and whether the process has resulted in significant energy transfer to the environment. Spectra of SDSS LRGs have revealed numerous examples of gravitationally lensed background galaxies, many of which form near-perfect Einstein rings (Bolton et al. 2006; the SLACS project). These systems provide powerful probes of the mass profiles of early-type galaxy halos (Koopmans et al., 2006), and BOSS will allow these studies to extend to greater lookback times with improved statistics, revealing the evolution of the dark matter potential wells of early-type galaxies in addition to their stellar components.

The QSO survey will provide a large sample of less luminous QSOs at $z = 2.5$, near the era of peak QSO activity, which (ironically) is the most poorly studied redshift range. This will constrain the faint end of the QSO luminosity function and provide the best data set for QSO clustering at these redshifts. The dependence of the clustering on QSO properties, such as luminosity, provides a novel constraint on the lifetimes and hosts of these rare objects and can be used to discriminate among the array of currently popular models for AGN feedback in massive galaxies. Surveys of high-redshift QSOs have been one of the high impact areas for SDSS-I. These objects trace the evolution of early generations of supermassive black holes, provide tests for QSO formation and AGN evolution, and probe IGM evolution. Due to the high impact of the highest redshift quasars and the relatively low cost of obtaining spectra for these rare objects, we will target higher z quasars as a piggyback program. By using only a small number of "extra" fibers BOSS should approximately double the number of $z > 3.6$ QSOs discovered in SDSS and reach about one magnitude fainter in the luminosity function. QSOs at $z > 6$ display complete Gunn-Peterson troughs, suggesting that the reionization phase transition completed just shortly before $z \sim 6$. BOSS should roughly double the number of rare, $z > 6$ QSOs relative to the SDSS.

SDSS-I provided a great deal of stellar and Galactic structure science via calibration stars,

stars originally targeted as quasar candidates, and stars targeted for stellar science programs (see §4). We expect the same to be true for BOSS, with the bonus of higher sensitivity and somewhat higher spectral resolution. Since SEGUE-2 will already be carrying out a major optical stellar spectroscopic survey, we do not plan a large stellar science program in parallel with BOSS, but BOSS can easily accommodate a survey of rare objects (stellar or non-stellar), up to a few per square degree.

To recap, BOSS is a high-precision experiment to measure the acceleration of the Universe using the baryon acoustic oscillation technique, which has a low level of systematics. The proposed survey builds on the highly successful SDSS, utilizing the skills developed by the SDSS team and a significant fraction of the SDSS infrastructure. The BAO measurements from BOSS will only increase in scientific value as other dark energy projects progress, including other BAO surveys that focus on higher redshifts. This is true even if the currently "conventional" descriptions of dark energy as a smooth, slowly evolving field prove to be correct, since the BOSS anchor of the low-redshift distance scale and the precise measurement of space curvature from the BOSS quasar survey both provide complementary information that increases the leverage of other measurements. It is even more true if the origin of cosmic acceleration is (still) more exotic, as in models that appeal to modified gravity or extra dimensions instead of a new energy component. To reveal the signatures of such models, it is essential to combine techniques with complementary information, in particular a purely geometrical method like BAO with a method (such as weak gravitational lensing or the abundance of galaxy clusters) that measures the gravitational growth of structure. The "guaranteed" return of BOSS is much tighter constraints on the parameters of dark energy, but the "discovery" potential is larger still.

4 SEGUE-2: Mining the Outer Regions of the Milky Way

4.1 Introduction

Recent detailed studies of old stellar populations in the Milky Way and nearby galaxies have revolutionized the way we look at galaxy formation by directly probing the history of our Galaxy and its neighbors, providing complementary information to that from observations of the high-redshift Universe. The impact of these "near-field cosmology" investigations on our understanding of galaxy formation is due to the extensive kinematic and density information provided by massive surveys like SDSS and SEGUE. The old stellar populations in the Galaxy's thin disk, thick disk, halo, and bulge were born during the most active period of its formation, when models for galaxy formation predict the Milky Way gained most of its mass through hierarchical merging and accretion. The evidence for how those populations became part of the Galaxy, and for the history of element formation, is encoded in the kinematics, spatial distribution and chemical abundances of the stars and will be probed by SEGUE-2 (and, subsequently, by APOGEE).

Large-area sky surveys like 2MASS, QUEST and SDSS show that the outer Galaxy is marked by an abundance of substructure, signatures of the hierarchical assembly of the Milky Way from smaller systems (see Figure 7). Tidal tails and substructure from merger events completed or in progress have been found in the stellar density distribution in the Galactic halo (Newberg et al. 2002; Yanny et al. 2003; Crane et al. 2003; Ibata et al. 2003; Majewski et al. 2003; Belokurov et al. 2006a), including the recently discovered "Virgo Overdensity" (Vivas & Zinn 2006; Jurić et al. 2007; Bell et al. 2007; Martínez-Delgado et al. 2007), which covers so much sky that only a large-area survey like SDSS can isolate it from the foreground disk stars and reveal its vast extent.

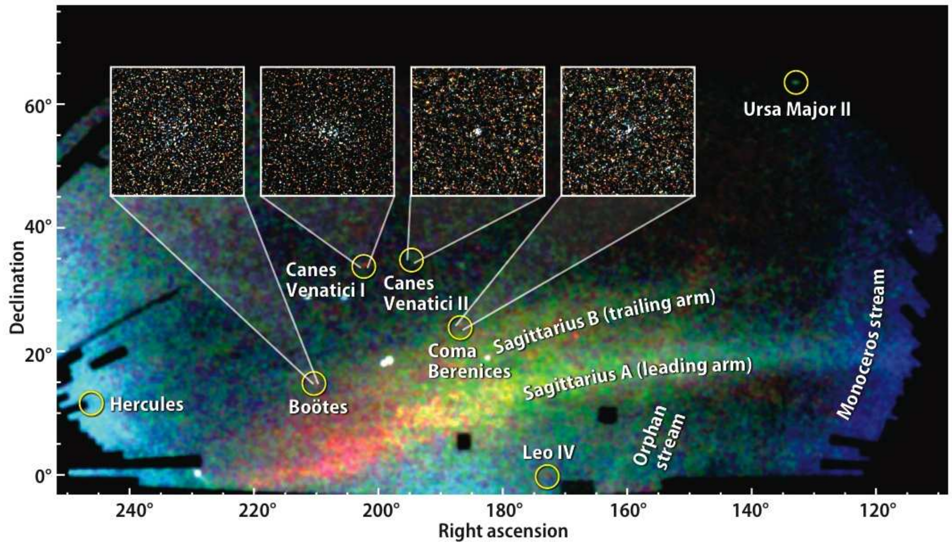

Figure 7: A map of stars in the outer regions of the Milky Way, covering about 1/5 of the sky, as observed by the SDSS. The trails and streams that cross the image are stars torn from disrupted Milky Way satellites (Belokurov et al. 2006a). The inset panels show four of the surviving, but very faint, dwarf galaxies discovered by the SDSS. The color-coding represents distance, with red being the most distant and blue being the closest; the distances of stars are inferred from their colors and magnitudes in the 5-band SDSS images.

SDSS has also led to the discovery nine new, very faint satellite galaxies that appear to be distinct substructures orbiting in the Milky Way's dark matter halo (Willman et al. 2005a,b; Zucker et al. 2006a,b; Belokurov et al. 2006a,b, 2007; Irwin et al. 2007).

As the velocity data from surveys like SEGUE become available, it has become clear that halo velocity fields are complex. The increase in sample size in SEGUE-2 and its emphasis on sampling the distant halo will shift the focus from the discovery of substructure to quantitative descriptions of the kinematics and chemical abundance distributions of stars in the halo and outer Galaxy. Figure 8 shows a comparison of velocity histograms from SEGUE data with the Law et al. (2005) model of the Sagittarius tidal stream. The dense sampling and increased number of fields in SEGUE-2 will provide better constraints on the orbits of streams like Sgr, as well as the numbers required to detect lower-density substructure left from earlier merger events. The *combination* of velocity and chemical abundances greatly enhances the ability to identify and characterize complex structure. For example, Carollo et al.'s (2007) analysis of 20,000 SDSS-I calibration stars shows that *inner halo* stars have a somewhat flattened spatial distribution, elongated orbits that on average follow the sense of rotation of the disk, and a metallicity distribution peaking at [Fe/H]$= -1.6$, while *outer halo* stars have a spherical distribution, retrograde mean rotation, and a peak metallicity [Fe/H]$= -2.2$.

The emerging picture of the lumpy Milky Way is in qualitative agreement with current theoretical expectations. Simulations of galaxy formation in individual dark matter halos based on initial conditions set by cosmological models predict the properties of Milky Way-sized dark matter halos and the galaxies they host in an average sense. The area and distance range of the SEGUE-2 data will build up a statistical description of the Galaxy over a large volume, providing the data needed

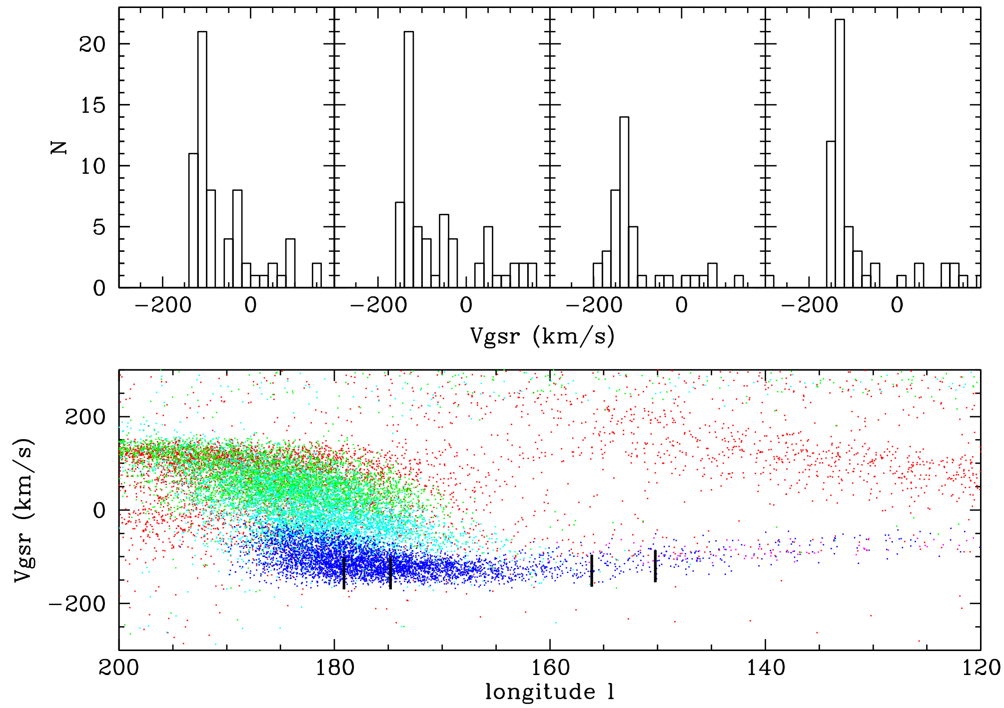

Figure 8: Observations and simulations of the Sgr stream. *(Lower panel)*: Model Galactocentric radial velocities vs. Galactic longitude along the Sgr stream, from Law et al. (2005). The points are color-coded according to the formation time of the stream. Blue = most recent passage, preceded by cyan, red, green. *(Upper panel)*: Observed Galactocentric radial velocity histograms for K giants from SEGUE at the locations indicated in the lower panel. There is good agreement between the data and the model for the most recent passage. The better sampling at larger distances in SEGUE-2 will improve constraints on earlier passages.

to constrain the models and inform their future development as they continue to increase in predictive power. The density profile and statistical distribution of substructure from high resolution simulations (Diemand et al. 2007) can be compared with SEGUE-2 data on the orbital properties of streams (Law et al. 2005; Fellhauer et al. 2006), orbits of individual stars in the inner halo (Carollo et al. 2007), and velocities of high-luminosity tracers of the outer halo (e.g. Battaglia et al. 2005; Dehnen et al. 2006). To these purely gravitational simulations are added hydrodynamic (Robertson et al. 2004; Brook et al. 2004; Stinson et al. 2006; Governato et al. 2007) and semi-analytic (Bullock et al. 2001, Bullock & Johnston 2005; Robertson et al. 2005; Tumlinson 2006) prescriptions for star formation, feedback, and the tidal accretion of satellites (Kravtsov et al. 2004; Bullock & Johnston 2005; Abadi et al. 2006; Moore et al. 2006; Read et al. 2006; Bell et al. 2007; Sales et al. 2007) as well as chemical evolution codes to model the evolution of element abundance patterns (Robertson et al. 2005; Font et al. 2006a,b). The chemical abundance information from the SEGUE-2 data, in combination with the kinematics, will be an important constraint on how the physical processes governing star formation in low-mass dark matter halos are implemented in these models.

Turning to the Galactic disk, there is now significant evidence for the existence of substructure in the old thin and thick disk components of our Galaxy (Eggen 1996; Gilmore, Wyse & Norris 2002; Navarro et al. 2004; Wyse et al. 2006; Helmi et al. 2006), suggesting that the same energetic accretion processes that built up the halo of the Galaxy played a major part in the formation of its disk. This type of complex disk assembly is predicted by recent, high-resolution hydrodynamic simulations (Abadi et al. 2003a,b; Brook et al. (2004, 2005). The SEGUE data include a color-selected sample that is dominated by *in situ* thick disk stars, and we will concentrate on completing the radial and Galactic longitude distribution of that sample in SEGUE-2. The resulting thick disk sample will be complementary to recent nearby, high-resolution spectroscopic studies of the thick disk (Fuhrmann 1998; Bensby et al. 2005; Brewer & Carney 2006) and will allow us to address the question of how important hierarchical merging and accretion were in the formation of the Galactic disk.

An equally important element of SEGUE and SEGUE-2 is the search for very low metallicity stars in the Milky Way, which provide fundamental constraints on the nucleosynthetic history of the early Galaxy and the lower-mass systems it accreted. Already, spectra from SDSS-I and SEGUE have roughly tripled the number of stars identified in the Galaxy with [Fe/H] < -2.0, and the selection criteria for the candidate metal-poor stars allow an unbiased view of the Metallicity Distribution Function (MDF) of the halo populations below [Fe/H]$= -2.0$. The MDFs of the inner and outer halo populations have already been shown by Carollo et al. (2007) to differ substantially, thereby making the connection between the chemical evolution and eventual accretion of the sub-systems from which the halo components of the Milky Way were formed. The rarity of the lowest metallicity stars in the Galaxy (Fe/H] ≤ -4.0) requires that a considerable effort be made to obtain sufficient representation of the most extreme populations with SEGUE-2.

4.2 Survey Strategy

The 2008-9 observing season will be devoted to SEGUE-2, doubling the number of SEGUE stellar spectra and reaching deeper into the halo. SEGUE-2 targets will be selected from the high-latitude imaging of SDSS-I and SDSS-II and from the additional 3500 square degrees of lower latitude SEGUE imaging undertaken in SDSS-II. One of the primary motivations for SEGUE-2 is to increase the density with which we sample the populations in the outer Galaxy. We will increase the number of distant halo tracers in the SEGUE-2 sample by re-balancing how fibers are allocated as a function of magnitude for some of the target classes, since the more nearby volume of the Galaxy is already densely sampled by the SEGUE pointings. We will also take advantage of our experience with the SEGUE data to optimize our efficiency for selecting distant halo tracers like red giants. The combined gain should be roughly a factor of four more stars in the distant halo. Otherwise, the targeting strategy will be closely based on that for SEGUE. SEGUE spectroscopy is limited to $|b| > 8^\circ$ because of limitations of crowded field photometry with the automated SDSS image processing pipelines. Since APOGEE is ideally suited to the study of the low latitude Galaxy, SEGUE-2 will concentrate on fields at $|b| > 20^\circ$ and at the Galactic anticenter.

Like SEGUE, SEGUE-2 will use color, magnitude, and proper motion to select stars that sample distant populations in the thick disk and halo, or other targets of special interest. Half of the ~ 1180 targets on each seven square degree line of sight are selected using straightforward color-color and color-magnitude cuts designed to sample across the main sequence from $g - r = 0.75$ (K-dwarfs at $T_{\rm eff} \sim 4800$ K) through the main sequence turnoff, weighted toward bluer stars with higher intrinsic luminosity on the main sequence. To this sample we add metal-poor main sequence turnoff stars selected by their blue *ugr* colors, essentially an ultraviolet excess cut that is extremely efficient at separating the halo from the thick disk near the turnoff. At r=19.5, the average star that makes this selection is at a heliocentric distance of 10 kpc for [Fe/H]= -1.54.

To reach greater distances, we use the strength of the Balmer jump to select field blue horizontal branch (BHB) stars in the *ugr* color-color diagram (Sirko et al. 2004; Clewley et al. 2004). The halo BHB sample extends to distances of 48 kpc for S/N > 15, the limit for stellar parameter determinations in SDSS DR6. We select all available BHB candidates. In addition, SEGUE selects candidate red giant branch stars in the halo by their offset location in the *ugr* color-color diagram relative to foreground disk dwarfs, the result of their ultraviolet excess and weak MgH (Morrison et al. 2001, Helmi et al. 2003). This is augmented by a 3σ proper motion cut using a recalibrated version of the USNO-B catalog (Munn et al. 2004). Spectroscopic identification of true giants using the methodology of Morrison et al. (2003) has shown that the giant selection is roughly 50% efficient at g <17, the limit to which we can reliably identify giants in the spectra as of DR6. The halo giant sample identified in this way reaches distances of 40 kpc from the Sun.

SEGUE also selects small categories of rare but interesting objects. These include cool white dwarfs selected with the recalibrated USNO-B reduced proper motion diagram, which can be used to date the age of the Galactic disk (Gates et al. 2004; Harris et al. 2006, 2007), and high proper motion targets from the SUPERBLINK catalog (Lépine & Shara 2005), which have uncovered some of the most extreme M subdwarfs known (Lépine et al. 2007, in preparation) and aided in the calibration of their metallicity scale using common proper motion pairs. The faint SEGUE-2 spectroscopic data are the most likely to yield new hyper-velocity stars (HVS; see Brown et al. 2005, 2007; Kollmeier & Gould 2007; Janiv et al. 2007, in preparation), which probe the dynamics and stellar populations of the Galactic center where they originate. While the known HVS are B stars, SEGUE-2 will be able to constrain the old population HVS, at or below the halo main sequence turnoff.

Beginning in Fall 2009, the bright-time SEGUE-2 parallel program will use the BOSS spectrographs, taking advantage of the improved sensitivity and the smaller ($2''$) fibers. SEGUE-2 will have 300-500 fibers for the bright-time program, about half of which will be devoted to sky to ensure good sky subtraction, and the available observing time should yield 100,000 or more stellar spectra. The magnitude limit is $g = 17$ for $S/N \geq 20$, faint enough to enable a wide range of interesting Galactic science. At $g < 17$ in the moderately low Galactic latitude MARVELS fields, the SEGUE-2 spectroscopy of main sequence stars can reach distances of 2.5 kpc, far enough from the plane to be dominated by the thick disk at most latitudes. In the absence of extinction, red clump stars and red giants are brighter than this limit out to the distance of the Galactic center, and AGB stars reach further still. The MARVELS and APOGEE fields will not always have SDSS imaging available, but we will be able to implement the selection of many of the same kinds of targets as for the main SEGUE-2 sample, including main-sequence turnoff, metal-poor and high proper-motion stars, BHBs and white dwarfs, using 2MASS, GALEX and other surveys. There are sufficient targets to fill 250 target fibers for multiple visits to the same field, even from imaging catalogs such as 2MASS that have brighter magnitude limits than the SDSS.

During 2008-2009 and the subsequent parallel program, we will carry out a series of calibration observations in coordination with the APOGEE team. Our plan is to obtain overlapping observations of the same well-studied open and globular clusters, targeting the same stars where possible, as well as duplicate observations of some field stars. This will ensure that both surveys can be placed on the same abundance scales and that the velocity observations are consistent across the two datasets.

The scientific utility of the SEGUE and SEGUE-2 spectroscopic sample rests on our ability to extract useful information from our moderate resolution spectroscopic data. The SEGUE team

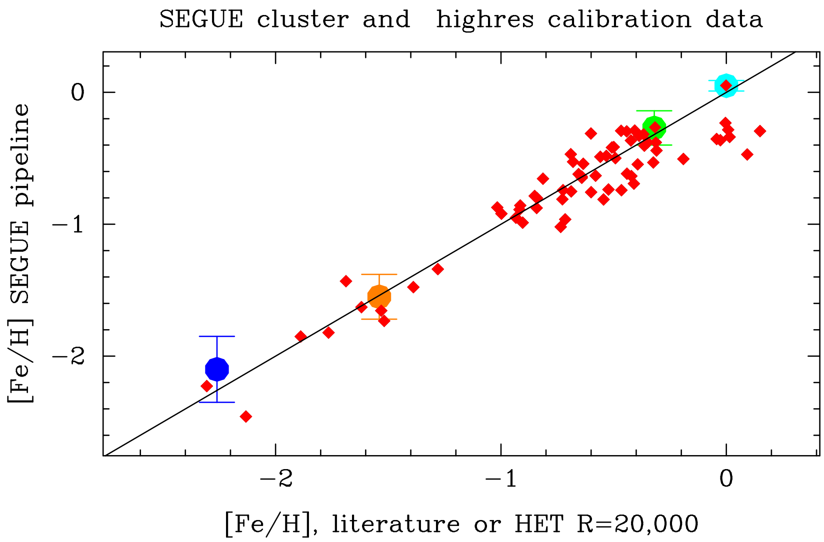

Figure 9: Comparison of known metallicity values for the SEGUE globular and open clusters (circles) and high resolution calibration sample from HET (red diamonds) with values measured by the SEGUE stellar parameters pipeline. The cluster color coding is: blue M15, orange M13, green NGC 2420, cyan M67. The vertical error bars plot the standard deviation of the SEGUE metallicity values around the mean for each cluster. The line is not a fit, but shows the 1-1 relation.

has developed a stellar parameters pipeline to measure effective temperatures, surface gravities, and [Fe/H] values from the SDSS-I and SEGUE spectra using a combination of methods previously described in the literature (Beers et al. 1999; Wilhelm et al. 1999, Allende Prieto et al. 2006; Re Fiorentin et al. 2007) and methods developed specifically for SEGUE. We have also obtained a large sample of calibration data to check those parameter determinations. We have acquired photometric and spectroscopic observations with the actual SDSS 2.5m system of well-studied open and globular clusters, with metallicities ranging from above solar to [Fe/H] $= -2.5$ (Lee et al. 2007, in preparation), and high-resolution spectra for 150 of the brightest SEGUE targets using the Hobby-Eberly, Keck and Subaru telescopes (Allende Prieto et al. 2007, in preparation). Figure 9 shows the current state of our analysis of this calibration sample, comparing the metallicity estimates from the SEGUE stellar parameters pipeline to literature values for clusters or to measurements from high resolution HET spectra. For data of good S/N ratio, the scatter between SEGUE metallicity estimates and those from high resolution spectra is ~ 0.15 dex.

We have tested the SEGUE radial velocity accuracy using repeat observations of stars in the DR6 sample and a noise model applied to high S/N SEGUE data for which we also have high resolution spectra. The two methods are in good agreement. The typical rms radial velocity uncertainty at the median S/N of the sample ($r \sim 18$), averaged over the full range in color, is $4.1\,\mathrm{km\,s^{-1}}$. For blue stars at the metal-poor main sequence turnoff, the radial velocity uncertainty at the median S/N is $5.5\,\mathrm{km\,s^{-1}}$; at $g-r = 0.5$ and S/N $= 10$ ($g = 19.8$) it is $9\,\mathrm{km\,s^{-1}}$. The global zero-point is calibrated with observations of open and globular clusters and the high-resolution sample. The resulting radial velocity precision is sufficient to resolve the $\sim 10\,\mathrm{km\,s^{-1}}$ velocity dispersions of dwarf spheroidal galaxies, which are likely also to be characteristic of the most recently accreted tidal substructure we expect to find preferentially in the distant halo. The SEGUE-2 velocities for brighter stars will be more precise than this, allowing us to use the local data to identify older substructure that has spread over its orbit and become colder (Helmi & White 1999).

The first substantial public release of stellar parameter data that meet the SEGUE science requirements was part of SDSS DR6 in June 2007. Work is ongoing to extend the range of the stellar parameter estimates in $T_{\rm eff}$, find optimal estimates at low signal-to-noise ratio, and measure additional quantities such as the abundances of carbon, α-elements, and other species.

4.3 SEGUE-2 In Context

The size of the SEGUE-2 sample, its focus on the old disk and halo populations in the outer Galaxy, and the resolution and wavelength coverage of the spectra make SEGUE-2 unique among current and near-term studies of the Galaxy, ensuring that it will remain an important archival resource as future projects and facilities come on-line. SEGUE has already proved capable of identifying substructure, candidate metal-poor stars, and other rare objects that are suited for follow-up on current 8-m class telescopes, and it is complementary to Galactic structure projects that make use of wide-field capabilities on larger telescopes to target fainter magnitudes at dense sampling. SEGUE-2 will provide an archived database of stellar population and kinematic information that can be used for planning targeted observations with the next generation of giant telescopes.

RAVE (Steinmetz et al. 2006) is a spectroscopic survey of bright stars ($9 < I < 12$) in the southern hemisphere, observing the region of the CaII triplet to obtain radial velocities accurate to better than 2 km/s. RAVE is optimized for finding local, old, dynamically cold substructure in the disk and nearby halo, exactly complementary to SEGUE-2's focus on the outer Galaxy.

The GAIA astrometric mission (Perryman et al. 2001; http://www.rssd.esa.int) is scheduled to launch at the end of 2011 and will obtain astrometric data for 10^9 stars brighter than $V = 20$, with an expected astrometric precision at the end of the mission of of 0.02 mas (1σ) at $V = 15$. These accuracies result in a 10% distance error at a distance of 10 kpc and a 1σ proper motion error of 1 km/s at 20 kpc. GAIA will also measure CaII triplet radial velocities accurate to 15 km/s at $V = 17$ and broad-band spectral energy distributions to $V = 20$. The SEGUE-2 archive will be a valuable resource for linking the bright GAIA spectroscopic sample and the full GAIA astrometric and photometric catalog to enable three-dimensional kinematic studies of the Galaxy to the full depth of the GAIA astrometric catalog. Matching the GAIA and SEGUE-2 catalogs will provide spectroscopically measured $T_{\rm eff}$, $\log g$, and [Fe/H] for 500,000 stars in the GAIA catalog to a depth at least two magnitudes fainter than the GAIA spectroscopy.

LAMOST is a dedicated 4000-fiber system that will conduct a Galactic survey (Zhao et al. 2006) of 5 million objects at a resolution of 1000-2000 with limiting magnitude V=20. A moderate resolution ($R = 5000 - 10,000$) follow-up survey will target candidate metal-poor stars to $B = 18.5$. The LAMOST team plans to use the SEGUE stellar parameters pipeline for their Galactic survey. SEGUE-2 will be completed before LAMOST gets underway and can inform its target selection and science planning.

WFMOS is a wide-field (1.5°) fiber spectroscopic survey instrument which, as currently envisioned, will be mounted on the Subaru 8.2m telescope and have $\sim$ 4500 fibers divided between moderate resolution and high resolution spectrographs. One of the project's defining scientific motivations is understanding the Galaxy's assembly history. The SEGUE-2 catalog is well-matched in magnitude to provide input targets for the high resolution mode for a "Chemical Tagging" (Freeman & Bland-Hawthorn 2002; De Silva et al. 2007) survey or high-precision kinematic follow-up of substructure in the outer Galaxy.

Sesar et al. (2007) have demonstrated robust separation of stellar populations in the large-area, homogeneous SDSS imaging data by using a sparse spectroscopic calibration of photometric metallicity and distance estimates. As future imaging surveys like VST, PanSTARRS and the LSST come on line, the deep, homogeneous SEGUE-2 spectroscopic sample will be invaluable as a calibration grid for photometric selection of stellar populations in the new surveys.

4.4 SEGUE-2 Science

SEGUE-2 is designed to obtain accurate kinematic, metallicity and other stellar parameter data over a large volume of the outer Galaxy, providing a simultaneous view of its structural, chemical and dynamical evolution. With these data, SEGUE-2 will address the following questions about the formation and evolution of the Galaxy and about its stellar populations:

Halo Substructure: The large increase in sample size in SEGUE-2 will give the statistics needed to characterize the frequency, filling factor and distribution of halo substructure with Galactocentric radius on a wide range of scales. With SEGUE-2's additional spectroscopy of outer halo fields we will have the sampling density to describe substructure in new dimensions (velocity and chemistry) as well as the three spatial ones provided by the photometry. This will allow comparisons with, and constraints on, models of galaxy formation tied to cosmological initial conditions and quantify the role of accretion in building up the stellar halo (Hartwick 1987; Preston et al. 1991; Zinn 1993a,b; Clewley et al. 2005; Abadi et al. 2006; Bell et al. 2007; Carollo et al. 2007).

Tracing Streams: SEGUE-2 will have the density of observed radial velocities to trace streams identified in the photometric data or distinct chemical signatures over large angles on the sky. The increased sampling density will allow us to identify stars on different wraps and separate streams that overlap spatially. Orbits fit to the streams provide stringent constraints on the Galactic potential.

The Dark Matter Halo: Using multiple probes of the outer halo, SEGUE-2 will determine the shape of the Milky Way's dark halo and constrain its large-scale density distribution. The wide-area substructure survey will enable detailed modeling of many more halo stream orbits than previously possible. Efficient selection of red giant and BHB stars will provide a large sample of kinematic tracers of the mass distribution at large distances from the Galactic center, which can measure the shape of the dark halo (e.g., Helmi 2004; Law et al. 2005; Battaglia et al. 2005; Dehnen et al. 2006). Increasing the sample of hyper-velocity stars, especially the old-population stars in the outer halo (Brown et al. 2005, 2007), will give information on the dynamical and star-forming conditions in the inner Galaxy and provide stars that are unique probes of the Galactic potential (Gnedin et al. 2005; Kollmeier & Gould 2007).

The Thick Disk: SEGUE-2 velocity data will sample the thick disk along many sightlines and in Galactocentric radius out to 15 kpc and quantify the frequency and radial distribution of substructure. The abundance data can also be used to look for differences in the stellar population parameters of the accreted component of the thick disk vs. the thick disk stars formed *in situ.* This census of substructure will allow us to test the idea that merging and accretion were important for the formation of the thick disk. With the combined unbiased kinematic and chemical abundance data in the SEGUE-2 thick disk sample we will be able to place quantitative limits on the number of significant merger events, the fraction of thick disk stars acquired through accretion, and the role of gas-rich mergers, and we will have the information necessary to disentangle these structures from perturbations caused by the Galactic bar (Dehnen 1998; Helmi et al. 2006)

The Metallicity Distribution Function: SEGUE-2 will measure the shape of the MDF at $[Fe/H] < -2$ with more than a factor of ten more stars than have ever been identified in this metallicity range. This robust measurement of the low end of the MDF will provide strong constraints on the chemical evolution of the Galaxy and of the progenitor structures that it accreted.

Chemically Primitive Stars: SEGUE-2 will identify candidate metal-poor stars at [Fe/H] < -4 for follow-up on large telescopes. Today there are only three stars known with [Fe/H]< -4.0 (Norris et al. 2007; Christlieb et al. 2002; Frebel et al. 2005). If SEGUE and SEGUE-2 are successful in even doubling this number, major new insight will be obtained into the elemental yields of some of the very first stars to have formed in the Galaxy, and indeed in the Universe. Such targets would be among the most valuable objects for detailed studies with present and future large aperture telescopes. The abundance patterns of these stars can constrain the initial mass function of Population III stars (Tumlinson 2006) and the yields of different nucleosynthetic processes. The outer-halo population possesses a MDF that extends to lower [Fe/H] than the inner-halo MDF; for example, all three of the known [Fe/H]< -4 stars are located at distances greater than 10 kpc from the Sun or have motions that carry them far into the outer halo. The focus of SEGUE-2 on the outer halo should help us find more such stars.

AGB Nucleosynthesis: The survey of neutron-capture elements in the SEGUE-2 parallel program will have AGB stars as a major class of target. The exact site of the s-process is still unknown, and the SEGUE-2 AGB star sample will allow us to isolate the mass, metallicity, and luminosity ranges of AGB stars that enrich the ISM with the s-process.

Neutron-Capture Elements in APOGEE Targets: The SEGUE-2 parallel program will measure the abundances of neutron-capture elements for the same Galactic populations as APOGEE. The AGB stars that are the site of the s-process are a major focus of APOGEE, and with the optical spectra SEGUE-2 can identify those AGB stars actively creating neutron-capture elements by the presence of optical YO and ZrO bands as they are dredged up at the tip of the AGB. At the distance of the Galactic center, these stars have H band magnitude of $\sim$8 and intrinsic (unreddened) g magnitudes of $\sim$13, easily observable by APOGEE and brighter than the $g = 17$ limit for the optical parallel data, so we can observe the same stars with both surveys. For populations without bright AGB stars, the SEGUE-2 parallel observations can target luminous tracers like red giants. By adding the abundances of neutron-capture elements to the APOGEE survey of the disk and bulge, SEGUE-2 will determine how different populations enriched the Galaxy in these elements.

Rare Objects: Both SEGUE-2 and the parallel/bright time program will produce sizable samples of rare objects. These will include samples of rare, chemically peculiar white dwarfs (WD) large enough to make informed analyses of their place in stellar evolution (Liebert et al. 2003; Köster & Knist 2006; Dufour et al. 2007), and samples of very cool WDs large enough to pin down the age and formation history of the local Galactic disk (Gates et al. 2004; Kawka & Vennes 2006; Kilic et al. 2006; Harris et al. 2007). Many DZ/DAZ WDs may owe their peculiar abundances to accretion of objects from the remains of their planetary systems, and circumstellar dust and gas disks are being found around an increasing number of WDs, in some cases using SDSS spectroscopy (Becklin et al. 2005; Kilic et al. 2005; Gänsicke et al. 2006, 2007; von Hippel et al. 2007; Jura et al. 2007). SEGUE-2 will continue to find very cool red dwarfs, which trace the bottom of the hydrogen-burning main sequence and provide targets for high-contrast, high-resolution searches for very low mass dwarfs and high-mass planets, allowing the investigation of planetary system formation. SEGUE-2 will also continue to find dwarf carbon stars and measure their peculiar abundances, which are presumably due to nucleosynthesis in a long-dead companion. Studies of large samples of these and other rare objects add information on the substellar initial mass function, the formation history of the Galactic disk, the production of elements in the early Universe, and the history of planetary systems.

5 APOGEE: Unveiling the History of the Milky Way

Much of the Milky Way Galaxy—particularly its inner disk, bar, and bulge— is obscured at visual wavelengths by interstellar dust (Fig. 10), which makes it nearly inaccessible to optical observations even with the largest telescopes. The Apache Point Observatory Galactic Evolution Experiment (APOGEE) will revolutionize our understanding of the evolutionary history of our Galaxy by surveying the chemistry and kinematics of over 10^5 stars—many of them located in regions never explored by optical studies. On the basis of high resolution and high signal-to-noise spectra of such a huge stellar sample, it will be possible to discern the true three-dimensional view of our Galaxy and read the story of how enrichment proceeded as a function of time and location, a story that is recorded in the chemical abundance patterns of its stars.

Optical studies of stars in the inner parts of the Galaxy are severely limited in depth, except for very few lines of sight that are "windows" in the screen of dust. Infrared observations, such as those provided by the Two Micron All-Sky Survey (2MASS, Skrutskie et al. 2006), are less hindered by dust extinction, and have allowed a much clearer view of the Galaxy (Fig. 10). Evolved stars, such as red giant clump, red giant branch (RGB), and asymptotic giant branch (AGB) stars, are readily identifiable in the 2MASS point source catalog, extend across the disk, bulge and halo, and, with the addition of data obtainable with other infrared facilities, provide a means to explore the structure, dynamics, and chemistry of the Galaxy in a uniform and systematic way. SDSS-I and SDSS-II have clearly demonstrated the enormous advantages of a homogenous and systematic survey strategy for bringing into sharp relief the detailed but subtle substructure of the Galactic halo. By contrast, studies of the structure, dynamics, and chemical evolution of the inner Galaxy have employed a heterogeneous mix of tracers — usually the easiest to see in each Galactic regime (e.g., planetary nebulae, HII regions, OB stars, Cepheid variables, open clusters) — and have used different techniques to explore them (e.g., different chemical species to test Galactic chemical evolution, emission versus absorption line studies, optical versus infrared). The uncertainties inherent in the comparisons of such disparate types of data are often larger than the astrophysical trends being sought. As a result, we still do not have a good understanding of the chemical trends across the Galactic disk, how they relate to the dynamics and orbits of stars, and how structures like the bulge, the bar, the thin disk, the thick disk, and the halo are connected to or differ from one another dynamically, chemically, structurally, and in terms of relative ages and star formation histories. Without these key constraints we cannot hope to formulate an integrated picture of the formation and evolution of our Galaxy, which is the best laboratory we have for studying the detailed evolution of galaxies in general.

APOGEE will exploit the ability of infrared observations to penetrate the veil of interstellar dust, performing the first large scale, systematic and homogeneous survey of the stellar content of our *entire* Galaxy. APOGEE builds on the success of the 2MASS photometric survey and is tied to this heritage through the APOGEE Instrument Scientist Michael Skrutskie, who was the Principal Investigator of 2MASS. APOGEE complements and extends the SEGUE and SEGUE-2 surveys in three critical ways:

(1) A Spectroscopic Survey of ***All*** *Galactic Stellar Populations:* With a 300-fiber, cryogenic spectrograph operating in the near-infrared (NIR) H-band (1.6μm), where extinction is much lower than in the optical (A_H=$A_V/6$, which translates to a factor of 100 in H-band to V-band flux at A_H=1), APOGEE will accurately determine the chemical content and kinematics of $\gtrsim 10^5$ stars, characterizing with unprecedented detail *all* stellar populations of the Galaxy, including those normally avoided in optical studies because of extinction. Unlike previous SDSS surveys, which had to create their own photometric databases before proceeding with the spectroscopic phase of the

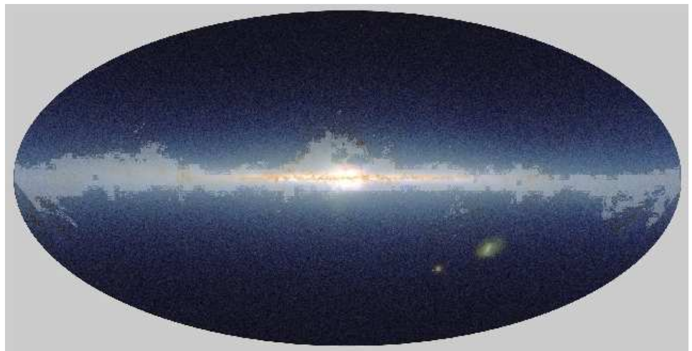

Figure 10: Panoramic view of the sky as seen by the Two Micron All-Sky Survey (2MASS), prominently showing the disk and bulge of the Milky Way. The gray overlay shows regions in which extinction at visual wavelengths exceeds 1 magnitude (a factor of 2.5 in flux), demonstrating the need for infrared spectroscopy to probe the Galactic disk.

survey, the corresponding photometric database already exists for APOGEE, which will uniformly select giant stars (RGB, AGB, red clump) to a flux limit of $H\sim13.5$ directly from 2MASS. This focus on one particular type of tracer that is both intrinsically luminous and found in stellar populations of all ages further ensures the uniformity of the APOGEE survey across all parts of the Galaxy.

(2) High-Resolution Chemical Abundances: In a revolutionary advance, APOGEE will obtain spectra at *high resolution* ($R = \lambda/\Delta\lambda \sim 20,000$) and *high S/N ratio* ($\sim$100 per resolution element) for these $\gtrsim 10^5$ giant stars. (For comparison, SEGUE has $\lambda/\Delta\lambda \sim 2000$ and typical S/N ratio ~ 25; high-resolution spectroscopy of SEGUE stars requires expensive follow-up observations with large telescopes.) With its high resolution and S/N ratio, APOGEE will go beyond measuring [Fe/H] and [α/Fe] and will instead determine for each star detailed, accurate *abundance patterns* spanning numerous chemical species and providing unprecedented sensitivity to different nucleosynthetic pathways. These measurements will yield far more detailed insights into the chemical evolution of the Galaxy than can be construed from [Fe/H] and [α/Fe] alone.

APOGEE will measure the abundances of nuclei that are produced by distinct physical processes that operate in stars of different masses, metallicities, and lifetimes, and which therefore probe distinct environments and/or distinct epochs of Galactic history: (1) The H-band contains numerous lines of the OH, CN, and CO molecules, which can be used to determine abundances of oxygen (from OH), carbon (from CO, but also from C I), and nitrogen (from CN). C, N and O are the most abundant heavy elements in the Universe and are produced prodigiously in the SN II that were the primary method of enrichment at early times; thus, these are the preferred reference species for gauging the overall chemical evolution of a stellar population. (2) The H-band also has lines from almost the entire sequence (O, Mg, Si, S, Ca, Ti) of α-elements — those nuclei that are multiples of the helium nucleus and which are also produced primarily in the explosions of short-lived, massive stars. The relative abundances of these elements reflect the shape of the stellar initial mass function (IMF) contributing to the enrichment of the interstellar medium. (3) Elements in the "iron peak", including Cr, V, Mn Fe, Co and Ni, also have H-band lines. These elements are produced copiously in the Type Ia supernovae of longer lived, lower mass stars, and have yields that are metallicity-dependent. (4) Finally, APOGEE will be able to sample elements with odd numbers of protons (the "odd-Z" elements Na, K and Al), which are produced mainly by massive stars (as are O and Mg) but whose yields from SN II, unlike O or Mg, are sensitive to *both* the IMF and the overall heavy-element content of the input star-forming gas. Thus, these elements tend to be strongly enhanced in systems where star formation proceeded at a slow pace.

APOGEE's rich trove of abundance data for all of these key elements in Galactic stars will provide detailed information on the star formation history, chemical enrichment pathways, and distributions of initial stellar masses within the parent stellar populations in the Galaxy. Some of the above elements (such as Mg, Ca, Ti, or Fe) are well-studied in the solar neighborhood via optical data, while others (such as S, K, Co) have not been studied as extensively or are generally unavailable in optical stellar spectra (as is the case for the all-important reference element oxygen). This kind of valuable and detailed information on the chemistry of Galactic stars is currently available for only a few hundred stars beyond the immediate solar neighborhood. APOGEE will increase this sample by 2–3 orders of magnitude, shedding new light on previously only minimally explored parts of the Galaxy.

(3) Precision Radial Velocities: APOGEE will provide the most comprehensive sample of precise radial velocities across the Galactic bulge, bar and disk, allowing the dynamics of these systems to be explored in detail. With accurate spectroscopic parallaxes to disk tracers, the Galactic rotation curve can be measured, even outside the solar circle where traditional H I tangent point analyses cannot be applied. The derived velocity field can be used to search for second order effects, e.g., the influence of the Galactic bar and spiral arms on the disk. Current models for the formation of galaxies in the ΛCDM cosmology imply that the assembly of the Milky Way was a complex, hierarchical process, punctuated by frequent merger events (e.g., Gilmore et al. 2002) with associated star formation and dynamical perturbations. As already discussed in §4, the stellar halo of the Milky Way bears clear imprints of this complex history, but we presently have little data with which to test such models in the inner Galaxy. APOGEE will provide stellar radial velocities with better than 0.5 km s^{-1} accuracy, allowing the *dynamical granularity* of the Galaxy to be accurately measured. Because collisionless evolution conserves phase space density, accreted structures that become spatially mixed simultaneously sharpen their contrast in velocity space, which may make them *easier* to pick out from the background (e.g., Helmi & White 1999, Helmi et al. 2006). APOGEE will resolve these residual phase space structures, and its ability to identify distinct, accreted populations will be dramatically enhanced by "chemical fingerprinting" of stars that bear the imprint of common formation history in their detailed abundance patterns. With the combination of precision velocities and abundance patterns, APOGEE will be the first survey capable of testing models of galaxy formation not only in the relative calm of the halo but throughout the bulge and disk of our Galaxy.

APOGEE will be the first large-scale, high-resolution spectroscopic survey spanning *all parts* of the Milky Way. Even in fields with 15 magnitudes of visual extinction, giant stars sampled by APOGEE will probe to 25 kpc, allowing a thorough survey across virtually the entire disk, bar, and bulge. The resulting uniform database of very high S/N spectra of Galactic stars will surpass by more than two orders of magnitude the total number of all high resolution, high S/N stellar spectra ever taken on all of the world's telescopes.

5.1 The APOGEE Spectrograph

APOGEE will extend the original SDSS fiber spectroscopy concept to infrared wavelengths and substantially higher spectral resolution via an innovative, multi-object, $R \sim 20,000$ spectrograph operating in the NIR "H-band", covering 1.52 to 1.69 μm (Figures 11 and 12). As with the optical SDSS spectrographs, light from a pseudo-slit composed of 300 ultra-low OH ("dry") fibers

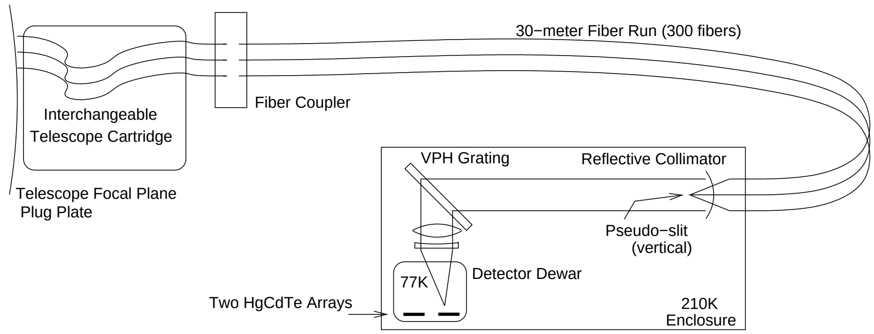

Figure 11: Schematic view of the APOGEE system from the 2.5-meter telescope focal plane to the detector arrays.

is collimated by an $f/4$ spherical mirror. At the telescope end APOGEE will use a plug-plate configuration and fiber harness similar to the original SDSS configuration (see §7). Each fiber core will subtend $2''$ on the sky corresponding to a 120μm core diameter (210μm overall fiber diameter). The collimator will relay light to a 385 mm diameter, 983 lines/mm, volume phase holographic (VPH) grating. A refractive $f/1.2$ camera then focuses the spectrum onto two 2048^2 HgCdTe detector arrays with 20 μm pixels. The arrays are staggered in the dispersion direction to provide optimal spectral coverage across the H-band. The arrays will be simultaneously clocked and read-out by a standard SDSU controller. Because thermal radiation is non-negligible at 1.6 μm, the spectrograph optics will be housed in a cryostat and held at a temperature <210K while the arrays operate near liquid nitrogen temperature. A blocking filter precedes the arrays to remove short wavelength spectral orders and to block thermal radiation at longer wavelengths. The spectrograph is too massive to ride on the telescope and will be floor-standing in an enclosure adjacent to the telescope in the expanded support building. The fiber pseudo-slit inside the dewar will connect to the telescope's focal-plane fiber cartridges via a 30-meter trunk of fiber terminating in 10 fiber couplers, each carrying approximately thirty fibers. These couplers plug into mating couplers on the focal-plane cartridge fiber harnesses. The couplers enable the exchange of the focal plane plug-plate cartridges during the night, since the fiber pseudo-slit is permanently mounted inside the cooled spectrograph.

5.2 Survey Strategy

APOGEE will acquire $S/N = 100$, $R \sim 20,000$ spectra of giant stars in multi-directional probes to $H = 13.5$, a brightness limit that represents a remarkable confluence of expected instrument sensitivity with several practical astronomical considerations. At this brightness limit and in the absence of extinction, red clump stars (the bluest and lowest luminosity of APOGEE's primary stellar tracers) lie at $H \sim 13.5$ at the distance of the Galactic center. An unreddened $H = 13.5$ star at the tip of the RGB lies at a distance of 70 kpc, and at 25 kpc when seen through 2.5 magnitudes of H-band extinction (which implies $A_V = 15$—a regime untouchable with optical probes). AGB stars are up to 1.5 mag brighter still. Finally, $H \sim 13.5$ represents the faint limit for reliable RGB/red clump identification within 2MASS.

Taking advantage of the lower mean atmospheric extinction and seeing degradation at high airmasses in the H-band compared to the optical, APOGEE can usefully probe the Galactic mid-plane from $l = -5^\circ$ to $l = -250^\circ$. Therefore, the bulge, bar, and inner and outer Galactic disk will

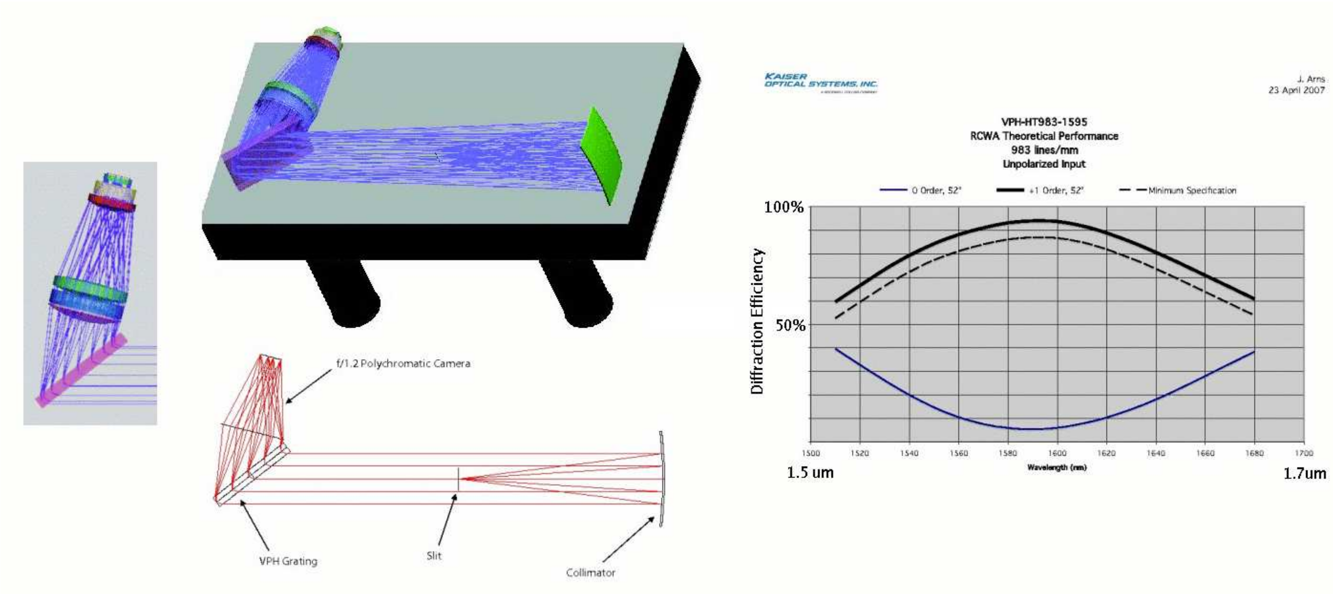

Figure 12: *Left*: APOGEE optical design and implementation on a 4'x8' optical bench. Inset at far left shows a proof-of-concept six-element optical design for the system camera developed within the SDSS-III project. *Right:* Vendor predictions of VPH efficiency peaking at 90% in middle of the spectral range and falling to 50% at the ends of the spectral coverage.

be explored, including fields of special interest such as Baade's Window, the Galactic Center, low-latitude globular and open clusters, and fields previously probed by SEGUE. Even considering the limits due to differential refraction, which impose restrictions on field size at higher airmasses, there will be no problem identifying targets even at the declinations characteristic of Baade's window and the Galactic center, given the high density of sources in these directions.

With conservative estimates on the instrument/telescope transmission (10%), interline airglow emission background (150 μJy), thermal emissivity of the fiber spectrograph (20%), thermal background, detector dark current (0.04 electrons sec^{-1}), and detector readnoise (8 electrons, after limited Fowler sampling), the sum of all noise equivalent and background photons/electrons is slightly smaller than the number of detected photons from an H=13.5 source, and leads to S/N= 100 in slightly more than three hours of total integration. Given a three year campaign with half of the bright time dedicated to APOGEE and typical weather, APOGEE could observe some 400 distinct fiber plate pointings with a total of three hours of integration each. With $\sim$ 250 of the 300 APOGEE fibers on 2MASS giant star targets (and the rest of the fibers on sky, to monitor airglow), a survey of at least 10^5 stars can be achieved.

In practice, we anticipate a total sample 1.5–2 times larger than this. First, in many directions, particularly towards the Galactic bulge, it will be desirable to undertake multiple pointings at the same position of the sky but probing different distance/magnitude ranges, so that in many cases there will be shorter length exposures to measure just the brighter stars. Second, the efficiency of both SDSS-III bright time programs will be boosted by using the same plugplates to feed the MARVELS and APOGEE spectrographs simultaneously. Because both of these surveys have a primary interest in observations along the Galactic plane, coordinated field selection is generally straightforward. In particular, the primary MARVELS criterion of $>$ 120 stars with $V \leq 12$ is

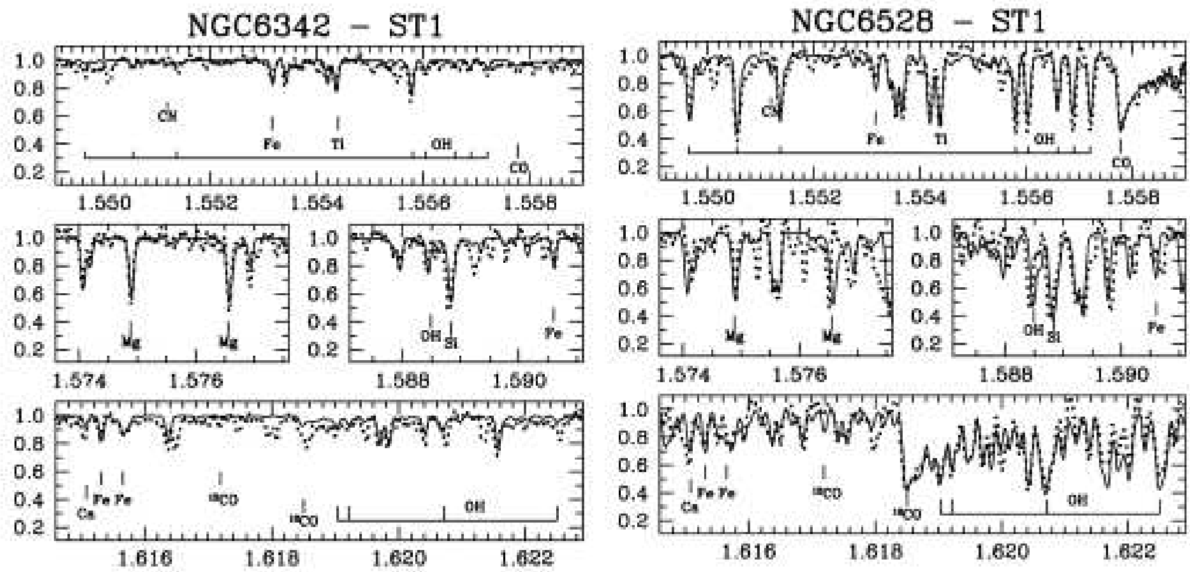

Figure 13: Sections of H-band spectra for two globular cluster giant stars showing additional α-element, Fe and CNO lines. These Keck+NIRSPEC spectra are at resolution $R \sim 25000$, similar to the $R \sim 20000$ resolution planned for APOGEE. Left spectrum is for an [Fe/H]$= -0.60$ red giant and right is for an [Fe/H]$= -0.17$ giant star. This is a typical metallicity range for Milky Way disk stars, and, as may be seen, this variation in abundances is obvious in these lines. The solid lines in the figures show spectral synthesis models being applied to interpret the abundances of the $R \sim 25000$ spectra (shown by the dotted lines). From Origlia et al. (2005).

satisfied by almost any APOGEE field. The MARVELS baseline strategy involves ~ 30 1-hour visits to each field compared to a nominal total exposure of three hours for APOGEE targets. However, the number of $H \leq 13.5$ stars in a typical APOGEE field greatly exceeds the number of fibers, so APOGEE can observe fields for longer than three hours and probe multiple distance bins at constant (l, b). Furthermore, MARVELS can obtain its high-cadence observations in some APOGEE fields before 2011, so that the number of observations required during overlap with APOGEE is much less than 30. While details of the survey integration are still being investigated, it seems clear that the effective share of bright time available to APOGEE and MARVELS will be *well over* 50% for each program.

5.3 APOGEE Data Reduction and Analysis Pipeline

Building on the tools developed for optical spectra in the SDSS, the data reduction pipeline will extract 1-D spectra from the 2-D focal plane images. Airglow is highly variable on minute timescales and over degree scales, and will be removed by scaled versions of the numerous APOGEE sky fibers distributed around the FOV. Telluric absorption will be dealt with via regular observations of hot rapidly rotating stars or of ratios of daylight exposures through different airmasses. Wavelength calibration can potentially employ airglow lines, but the SDSS arc lamp spectra can also be used. As with SEGUE, radial velocities will be derived by standard cross-correlation techniques against a library of templates for stars of different temperatures and metallicities.

APOGEE will employ well established techniques to determine elemental abundances from analysis of high resolution NIR spectra of giant stars (e.g., Meléndez et al. 2003; Rich & Origlia 2005; Cunha & Smith 2006). Nevertheless, the task of determining $T_{\rm eff}$, $\log g$, and over a dozen accurate elemental abundances for more than 10^5 stars is unprecedented, and so it needs to be met by a novel approach.

The APOGEE abundance analysis pipeline will implement classic methods in an automated routine in order to yield accurate, homogeneous, and easily reproducible results in a speedy fashion. State of the art model atmospheres and spectral synthesis (e.g., Fig. 13) based on complete and accurate line lists will be at the core of the analysis. Effective temperatures, surface gravities, and iron abundances will be determined using 2MASS photometry, theoretical isochrones, model atmosphere computations, and spectral synthesis of numerous Fe lines. The abundances of oxygen, carbon, and nitrogen will be determined from synthesis of rotational lines of the OH, CO, and CN molecules. The abundances of other elements such as Mg, Si, S, Ca, Ti, Cr, V, Mn, Ni, and Al will be based on synthesis of atomic lines.

The primary computational effort will actually take place *before* APOGEE observing begins, when a massive multidimensional library of many thousands of synthetic spectra will be developed, spanning the expected range of atmospheric parameters: $\log g$, $T_{\rm eff}$, ξ, [Fe/H], and [α/Fe]. The computational work before and during the survey will be reduced substantially by minimizing the size of the total spectral region synthesized via selection of sets of narrow spectral windows sensitive to the variation of each stellar parameter (e.g., Schiavon et al. 1997). The abundance reduction pipeline will proceed by extracting those windows from each APOGEE spectrum and searching the synthetic library for the best template match in these windows, based on standard χ^2-minimization routines, which combine χ^2-gradient search and linearization methods to locate best-matching sets of input parameters (e.g., Valenti & Fischer 2005). Individual elemental abundances beyond iron will then be determined by fitting the sets of relevant atomic/molecular lines with the stellar atmosphere fixed and the [X/H] varied. Based on previous analyses of H-band giant star spectra at similar resolution (but lower S/N; Rich & Origlia 2005), element ratios good to 0.1 dex or better in [X/Fe] are expected.

To achieve accurate absolute calibration of its abundance scale, APOGEE will target well-studied star clusters with a range of chemical compositions, as well as bright, "standard" stars, such as the Sun (from daylit sky), Arcturus, and μ Leo, and sources from the ELODIE library (Prugniel et al. 2007). These reference stars will overlap significantly with those being used to calibrate SEGUE/SEGUE-2 to ensure a robust consistency between all SDSS-III/SDSS stellar surveys. Furthermore, APOGEE can revisit a selected number of SEGUE fields, specifically cross-targeting stars already identified by SEGUE as particularly interesting, and creating a subsample of stars having not only detailed information on H-band accessible elements, but also a complement of neutron capture elements observable in the optical spectra.

Parameters included in APOGEE data releases will include, for each star, the pipeline-derived stellar atmosphere parameters ($\log g$, $T_{\rm eff}$, ξ, [Fe/H]), the elemental abundances for $\gtrsim$15 chemical elements, and the fully calibrated H-band spectra.

5.4 APOGEE in Context

(1) Spectroscopic Surveys: When contrasted with ongoing and future spectroscopic surveys of the stellar content of the Milky Way Galaxy, APOGEE stands out because of two of its defining properties: spectral resolution and depth. This is illustrated in Figure 14, where APOGEE is

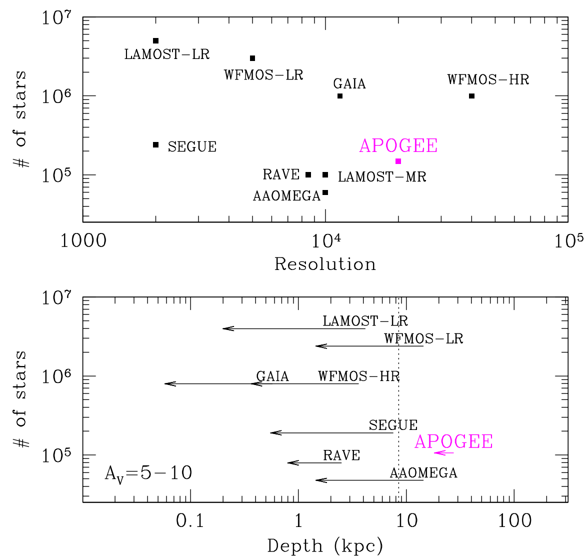

Figure 14: *Upper Panel:* APOGEE is compared with other ongoing and future surveys in the *Spectral Resolution* vs. *Sample Size* plane. Only APOGEE and WFMOS-HR have high enough resolution to perform a high quality detailed elemental abundance analysis. *Lower Panel:* Same comparison, now in the *Depth* vs. *Sample Size* plane. Depth here is defined as the distance at which a 1.2 $M_\odot$, 5 Gyr-old, solar metallicity giant would be observed at the limiting magnitude of each survey, adopting $A_V = 5$. The arrows indicate the change for $A_V = 10$. The dotted line marks the distance towards the Galactic center.

compared with the Milky Way spectroscopic surveys described in Section 4.3. The upper panel shows the relative spectral resolution versus the intended sample sizes for these surveys. Only two surveys — APOGEE and WFMOS-HR — have high enough resolution that unblended absorption lines can be used for high quality elemental abundance analysis. WFMOS-HR is projected to have twice the resolution of APOGEE, but this is largely compensated by the fact that absorption line crowding is far less important in the NIR than in the optical, which makes blending a less serious problem in the *H*-band. The current future of WFMOS is uncertain, but under any projection, APOGEE's scientific yield is expected to be delivered substantially earlier than WFMOS-HR. Apart from its more immediate impact, APOGEE can serve as an important source of interesting targets or directions of the sky for WFMOS-HR, should that project move forward.

The primary difference between the present proposal and the other Galactic spectroscopic surveys is that *all* of the other projects are in the *optical*. The APOGEE *H*-band survey brings a new, uniquely competitive and cost-efficient approach to the study of Galactic chemodynamics, particularly in studies of the highly dust-obscured parts of the Galaxy where optical spectroscopic programs would require the largest telescopes in the world to probe *some* of the lines of sight accessible to APOGEE. The lower panel of Figure 14 demonstrates the relative distance at which one would observe a 5 Gyr-old, 1.2 $M_\odot$, giant star with solar metallicity, at the limiting magnitude and characteristic photometric band of each survey, assuming $A_V = 5$. The arrows indicate how distances would change if $A_V = 10$ were adopted instead. Numerous Milky Way lines-of-sight near the midplane, and particularly towards the Galactic bulge, have fields with $A_V > 10$ and cannot be explored with RAVE, AAOmega, LAMOST, SEGUE or WFMOS. Only an infrared survey like APOGEE can provide a deep, high-resolution and comprehensive spectroscopic view of the inner Galaxy. Even high-resolution studies of Baade's Window, one of the most accessible portals into the heart of the Galaxy, have depended almost exclusively on recent NIR spectroscopy.

With its view of even the dust-enshrouded populations of our Galaxy, and its specific focus for all of its targeting on a single, photometrically identifiable stellar type — luminous giant stars — only APOGEE can both efficiently and uniformly explore all of the principal components of the Milky Way. In addition, because of its focus on the H-band, APOGEE delivers a different set of chemical species than the optical studies. A most important contrast is that optical studies cannot determine oxygen abundances — a serious limitation because oxygen is the most abundant of all metals and critical to understanding of the history of star formation for populations of all ages.

(2) Astrometric Surveys: Gaia will provide proper motions to the 30-50μas yr^{-1} level to $V \sim 20$. Because K5 giant targets have $V - H \sim 3.5$, for such APOGEE targets at $H = 13.5$ we can expect Gaia proper motions as long as $A_V \lesssim 3$. Note that the Gaia radial velocity program will *not* obtain velocities for stars at these magnitudes (nor will most RV surveys planned or underway). Because APOGEE is expected to be underway before NASA's SIM PlanetQuest, any interesting stars with $V < 20$ uncovered by APOGEE could also potentially be observed with this higher precision (few μas yr^{-1}) instrument. Before either of these space missions, existing Hipparcos, Tycho and UCAC proper motions can provide some proper motion data with sufficient precision to explore the dynamics of stars within a few kiloparsecs of the Sun.

(3) Photometric Surveys: A key aspect of the presently proposed survey is that *no new photometric observations are required*. Because the target list and all necessary photometry to interpret the spectroscopic data are provided by 2MASS, the survey is self-contained. However, APOGEE will have overlap with a number of other near and mid-infrared photometric surveys that have the potential to enhance the usefulness of the data. Large-area Spitzer/IRAC surveys (like GLIMPSE-I, II and III) already provide valuable 3.8μ and 5.6μ data in very obscured regions, which, when matched with 2MASS, allow for more precise dereddening of the 2MASS CMDs from which APOGEE targets are selected. NASA's Wide Field Infrared Explorer (WISE), scheduled for launch in late 2009, will provide an all-sky 3, 5, 12, and 23μm survey for regions outside the GLIMPSE regions. ESO's Visible and Infrared Survey Telescope for Astronomy (VISTA) will be mapping a large area of the southern sky in the Z, Y, J, H, K_s bands starting in 2007. It will provide much deeper NIR photometry than 2MASS, making possible special very deep APOGEE probes in selected directions of the sky.

5.5 APOGEE Science

Galactic Stellar Populations: APOGEE will change the face of Galactic chemical evolution studies by substantially broadening the radial distribution over which solid and statistically significant chemical abundance data exist. APOGEE will investigate the physical mechanisms that connect the bulge, thin disk, thick disk, and halo, and it will probe how, when, and why these populations differentiate structurally, chemically, and dynamically during the growth of galaxies. APOGEE will extend the work of SEGUE-2 by defining the metallicity distribution function for *all* stellar populations. More broadly, by completing the first systematic survey of the *3-D distribution functions* of numerous key chemical elements sensitive to different nucleosynthetic pathways and timescales, APOGEE will reveal the history of star formation throughout the bulge, disk and halo. Through homogeneous, unbiased sampling of the thin disk, thick disk, and bulge giant stars, and by providing accurate spectroscopic parallaxes, metallicities and radial velocities, APOGEE will become the premier data set for probing the correlations of position, kinematics, and chemistry throughout the Milky Way (even more so when GAIA and SIM proper motions become available).

Hierarchical Formation of the Inner Galaxy: In "concordance" ΛCDM models the inner Galaxy is formed "inside-out" by the continuous accretion of material with higher angular momentum (e.g., Mo et al. 1998, Abadi et al. 2003b), but the active merging history demonstrated by these simulations on large scales remains a challenge to reconcile with the observed properties of structures on disk galaxy scales. For example, the existence of very old thin disk stars near the Sun suggests that the Galactic disk has remained relatively unperturbed to the radius of the solar circle for much longer than is easily accommodated by the active merging phase expected. Moreover, the "angular momentum catastrophe" (Navarro & White 1994, Navarro & Steinmetz 2000) results in condensation of baryons into model disks that are too small (Abadi et al. 2003a). With its derivation of precision radial velocities and chemical abundance patterns, APOGEE will provide new constraints on these formation models by searching for the expected residue of the merger process via velocity substructure and chemical fingerprinting. In this way, APOGEE can determine the relative contributions of *in situ* star formation and accretion of previously formed stellar populations in the inner Galaxy.

Population III: APOGEE observations will extend the archaeological reach of SEGUE-2 by excavating the bulge for the oldest population of stars, predicted by hierarchical models of galaxy formation to lie in the heart of the Galaxy. If some metal-free "Population III" stars form with low enough mass to survive to the present day, then APOGEE will discover them. Otherwise, it will show what the nucleosynthetic products of Population III stars were, providing unique insight into their mass distribution (Tumlinson 2006).

Halo Substructure: The combination of precision kinematical and abundance measurements over its enormous stellar sample will make APOGEE a tremendous tool for identifying and probing halo substructures, particularly those at low latitudes like the Monoceros ring encircling the Galactic disk (Newberg et al. 2002, Majewski et al. 2003, Ibata et al. 2003, Rocha-Pinto et al. 2003). For already identified tidal streams, APOGEE will measure accurate velocity dispersions, which provide information on the mass of the disrupted parent satellite and place limits on dynamical heating of streams by lumps in the dark matter halo.

Galactic Dynamics: Until its mass profile is reliably established, we cannot confidently compare our own Galaxy—whose internal constitution is known with unique detail—with other spirals, for which we often have better global information (e.g., rotation curve, inferred mass distribution, velocity moments) but much less detailed information on stellar populations. Although like most other spiral galaxies, the Milky Way appears to have a nearly maximal disk (Sackett 1997, Gerhard 2002), it seems to fall off the Tully-Fisher relation (Flynn et al. 2006). With stellar distances fixed accurately by quality spectroscopic parallaxes, APOGEE can map the large scale dynamics of the bulge, bar, and disk, thereby probing the global distribution of light and dark matter via the first comprehensive determination of the rotation curve to the outermost reaches of the disk, where it is poorly known at present. With a radial velocity precision better than 0.5 km s^{-1}, APOGEE will be able to discern subtle effects such as the perturbations of orbits by spiral arms and the Galactic bar. These studies will have even more impact when combined with GAIA proper motions, creating a database of stars with very precise phase-space coordinates.

The Galactic Bulge: Spheroids, including bulges of spirals, make up 50-70% of the stellar mass in the local universe (Fukugita et al. 1998). In our Galaxy the mass of the bulge is estimated to be 10-25% that of the disk. Yet, when and how bulges formed are still among the most fundamental open questions in the field of galaxy evolution. The Milky Way is an ideal place to explore these questions. To date the total number of bulge stars explored spectroscopically at *any* resolution numbers in the few hundreds, whereas high-resolution studies span only dozens of bulge giants (Fulbright et al. 2006, Rich & Origlia 2005, Cunha & Smith 2006). Moreover, because almost all previous spectroscopy focused on a single line of sight, Baade's Window, little is known about abundance, age and dynamical gradients in the Galactic bulge. By collecting high-resolution spectroscopy of large and unbiased samples of stars across the bulge, APOGEE will completely revolutionize our understanding of the bulge by mapping velocity fields and abundance distributions and by providing information not only on relative ages (through chemical patterns) but also on absolute ages (by establishing good abundances and distances/reddenings for stars seen in deep color-magnitude diagrams). By providing the first "integrated spectroscopic view" of the Galactic bulge, APOGEE will help in the interpretation of the integrated spectra of extragalactic spheroids.

The Galactic Bar: Very little is known about the dynamics and chemistry of the Galactic bar, including whether it is a short-lived or long-lived phenomenon and whether and how it may excite spiral arm modes and act to mix stars and gas within the disk. APOGEE can make a significant contribution to the understanding of the bar through the first major dynamical and chemical survey of its stars. This will allow a basic characterization of the bar and its rotation speed. By determining whether there are clear and unique correlations between stellar position, kinematics and metallicity, APOGEE will establish whether the bar is short or long-lived. Possible effects of the bar on the dynamics of both the disk and nearby halo will also be studied. These effects may be far-reaching: the Sun lives just outside the Outer Lindblad Resonance of the bar, just where dynamics might be expected to be most perturbed (Olling & Merrifield 1998). The close alignment of the bar with respect to the Sun means that we can easily measure streaming motions along the bar with just radial velocities, allowing us to look for the hypothesized kinematical-chemical correlations.

Legacy Survey of Low-Latitude Star Clusters: APOGEE will carry out the largest systematic spectroscopic exploration of low-latitude globular and open clusters, establishing their distances with improved spectroscopic parallaxes and velocity membership censuses and thus yielding improved estimates of cluster ages from isochrone fitting. Star clusters provide critical absolute age calibrations for the relative chemical evolution sequences established by abundance patterns. The velocity data on the clusters will help establish their orbits lending, insight into their dynamical evolution.

Star Formation: With its enormous spectroscopic sample and access to both the α and odd-Z elements, APOGEE will constrain the shape of the IMF in each of the Galactic stellar populations, and, within each, as a function of radius and metallicity/age. The initial mass function is one of the key observable constraints on models of star formation.

Interstellar Extinction: By obtaining quality spectroscopic parallaxes for stars across the disk, APOGEE, combined with existing or expected optical, near- and mid-IR data (e.g., from WISE, GLIMPSE, or Herschel), can map the 3-D distribution of Galactic dust and constrain variations in the interstellar extinction law, which reflect variations in the chemical constitution of the interstellar medium.

Together, SEGUE, SEGUE-2, and APOGEE will provide the key data sets needed to understand our Galactic home in a fully cosmological context. Other cosmological observations have achieved tight constraints on the matter and energy contents of the universe and on the initial conditions for structure formation, but the physics of galaxy formation and the behavior of dark matter on sub-galactic scales remain topics of intense research and debate. Any successful theory of galaxy formation must account for the properties of the galaxy we know best, including its structure, the relative distributions of stars, gas, and dark matter, and the histories of star formation and chemical enrichment. By revealing the story of the Milky Way in unprecedented detail, the SDSS-III Galactic structure surveys will transform this vision of "near-field cosmology" into reality.

6 MARVELS: Revealing the Formation and Dynamical Evolution of Giant Planet Systems

One of the most remarkable astronomical developments of the last 15 years has been the discovery of an abundant population of extra-solar planets. Surveys to date have detected over 230 extrasolar planets (see Butler et al. 2006 and Udry et al. 2007), of which approximately 90% were discovered by detecting the host star's reflex motion from measurements of the star's radial velocity (RV). Extrasolar planets reveal an astonishing diversity of masses, semi-major axes and eccentricities, from the short period "hot Jupiters", to planets in very elongated orbits, to planetary systems with multiple Jupiter-mass planets, to the super-Earth-mass planets with orbital periods of a few days (Butler et al. 2004; McArthur et al. 2004; Santos et al. 2004; Rivera et al. 2005; Lovis et al. 2006; Udry et al. 2006). A total of 18 transiting planets have been discovered[3]; these systems enable the measurement of many physical properties (see Charbonneau et al. 2007 for a review). If any single statement captures the development of this field, it is that the observations have continually revealed unanticipated diversity of planetary systems.

The standard, core accretion scenario of giant planet formation (see Lissauer & Stevenson 2007) predicts that planets like Jupiter form in nearly circular orbits, with periods of several years or more. Growth is initiated by coalescence of icy bodies, which cannot survive close to the parent star; once the solid core reaches $5-10$ Earth masses its gravity is strong enough to rapidly accrete surrounding gas and grow to Jupiter-like masses. The two greatest surprises of extra-solar planetary discoveries have been that many giant planets have periods below one year, sometimes as short as one day, and that many of these planets are on highly eccentric rather than circular orbits. The first finding suggests that many giant planets "migrate" inward after their formation, because of dynamical interactions with their natal gas disks. Various explanations have been proposed for the high orbital eccentricities, typically invoking gravitational scattering of planets after formation.

The goal of the Multi-object APO Radial-Velocity Exoplanet Large-Area Survey (MARVELS) is to provide a large sample of short-to-intermediate period giant planets obtained from a large, well-characterized, and homogeneous survey of stars with known properties. This will be a critical data set for testing the emerging theoretical models of the formation, migration, and dynamical evolution of giant planet systems. The key technical innovation behind MARVELS is a multi-fiber instrument that combines a fixed-delay interferometer with moderate dispersion spectrographs. This approach allows simultaneous, high throughput, high velocity precision measurements of many objects using reasonable detector sizes. In terms of the range of planet masses and separations to which it is sensitive, MARVELS will complement ongoing and planned radial velocity, transit, microlensing, and astrometric surveys, which probe lower mass or longer period systems but will not yield a comparably large sample of dynamically evolved giant planets.

We are now conducting a pilot survey at APO using a 60-fiber prototype of this design, known as the Keck ET1. We will use an upgraded ET1 for the first two years of the MARVELS survey; during this time we will build a cloned instrument (ET2) so that the final four years of the survey

[3] See http://obswww.unige.ch/~pont/TRANSITS.htm for a list of known transiting planets along with properties and references.

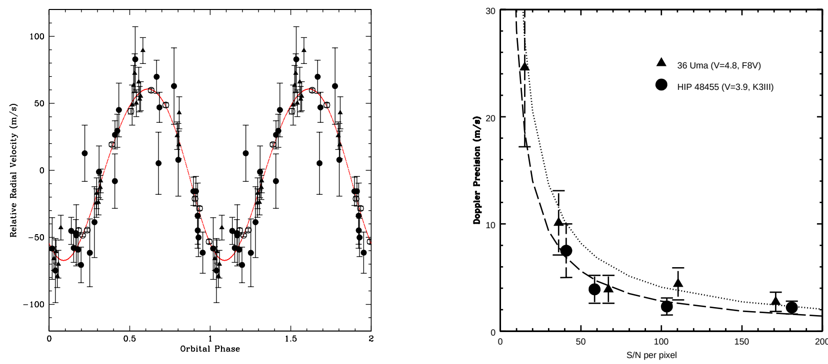

Figure 15: *Left:* The planet discovered with a prototype of the MARVELS spectrograph. Points show the measured radial velocity of the star against the orbital phase, with one repetition. The smooth curve shows the best-fit orbit, which is nearly circular with a period of 4.11 days. *Right:* Doppler precision measurements with KPNO ET in April 2007. The dotted and long dashed lines represents the photon noise limits for 36 UMa and HIP 48455, respectively.

will be performed with 120 fibers.

6.1 Survey Instrument Description

In the dispersed fixed-delay interferometer (DFDI) RV approach, a moderate-dispersion spectrograph is combined with a Michelson-type interferometer with a fixed delay between the two interferometer arms, adding a graded phase delay perpendicular to the dispersion direction which creates fringes in each resolution element. Doppler shifts cause the fringes in each resolution element to move (Erskine & Ge 2000; Ge 2002; Ge et al. 2002). This technique can employ relatively low resolution spectrographs with high throughput optics, and the data for a single star require only a small area in the detector plane; this latter property enables simultaneous RV measurements of many objects using a reasonably sized detector. The formula for the photon-limited Doppler precision of a DFDI instrument can be simplified to $\sigma_{\rm RV} \propto S^{-0.5}\Delta\lambda^{-0.5}R^{-0.5}D^{-1}$, where S is the stellar flux, $\Delta\lambda$ is the wavelength coverage, R is the spectrograph resolution, and D is the typical stellar absorption line depth. This formula resembles that for the echelle, $\sigma_{\rm RV} \propto S^{-0.5}\Delta\lambda^{-0.5}R^{-1.5}D^{-1}$ (Hatzes & Cochran, 1992); the main difference is the resolution dependence.

A DFDI instrument, the Exoplanet Tracker (ET), was commissioned at the KPNO 0.9-m Coude Feed/2.1-m telescope in late 2003. ET was used for a planet survey of $\sim$150 solar-type stars with $V = 7.6$–9 in 2004–2006; since fall 2006 the instrument has been available to the public. Figure 15 shows the short-term Doppler precision of the ET instrument — the dispersion of radial velocity measurements of the same star over a few hours baseline — as a function of signal-to-noise ratio. ET observations over two months of 51 Peg demonstrate that the instrument is stable to better than 5 m/s. ET has discovered one planet to date (Ge et al. 2006a; see Figure 15).

The first MARVELS survey instrument, constructed with support from the W.M. Keck Foundation and designated as the Keck ET, was commissioned at the APO 2.5-m in spring 2006. The Keck ET is based upon the design of the single-object KPNO ET. The Keck ET consists of eight subsystems: the multi-object fiber feed, the iodine cell, the fixed-delay interferometer system, the slit, the collimator, the grating, the camera, and the $4{\rm k} \times 4{\rm k}$ CCD. The instrument contains four

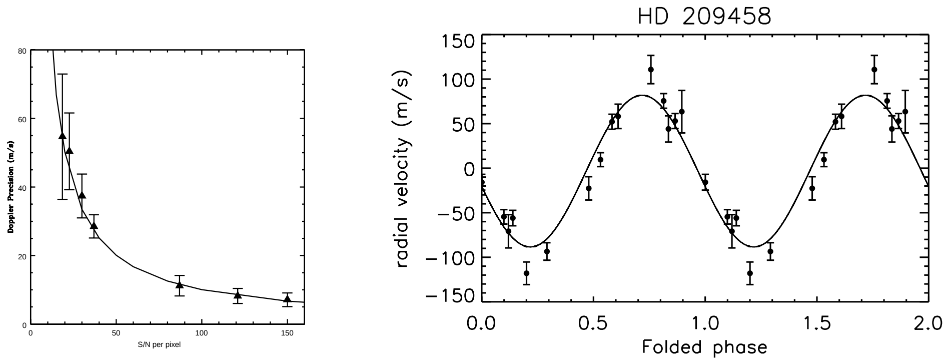

Figure 16: *Left:* Doppler precision measurements with Keck ET on 4 April 2007. The solid line is the photon noise limit. The triangles represent the 59-fiber average RMS error of the day sky RV measurements. *Right:* Keck ET observations of the previously known transiting planet system HD 209458. Points show the measurements (one orbital repetition); the smooth curve is the prediction based on earlier observations.

auxiliary subsystems for interferometer control, instrument calibration, photon flux monitoring, and thermal control. The instrument is fed with 60 fibers of 200 μm core diameters, which are coupled to 180 μm core-diameter short fibers from the telescope, corresponding to $3''$ on the sky at $f/5$. The spectral resolution for the spectrograph is R=5,100, and the wavelength coverage is 900 Å, centered at $\sim$ 5400 Å. Details of the instrument design can be found in Ge et al. (2006b), Wan et al. (2006), and Zhao & Ge (2006).

The current Keck ET has one spectrograph and one 4k $\times$ 4k CCD camera that capture one of the two interferometer outputs, and it has a 5.5% detection efficiency from the telescope to the detector without the iodine cell under the typical APO seeing conditions. The instrument can record 59 objects in a single exposure (a slight modification planned this fall will increase this to 60 spectra). The instrument Doppler precision was measured with the day sky scattered light, which offers a stable, homogeneous RV source for simultaneously calibrating the instrument performance for all of the sky fibers. The RMS error averaged over the 59 fibers, measured from the dispersion of measurements over a several hour interval in 2006 November, is 6.3$\pm$1.3 m/s. The corresponding average photon-limit error is 5.5$\pm$0.5 m/s (see Figure 16).

The instrument's precision over longer time intervals has been measured with repeat observations of sky scattered light over a period of 45 days in fall 2006, and 150 days in winter/spring 2007. The rms dispersion of RV measurements of sky over these periods, after subtracting the photon-noise errors in quadrature, are 11.7 $\pm$ 2.7 m/s and 11.3 $\pm$ 2.5 m/s, respectively. Recent studies, especially instrument performance simulations, show that the instrumental contributions to measurement errors are mainly caused by inhomogeneous illumination of the slit, image aberration, and the interferometer comb aliasing (sampling on the detector). Based on the current instrument performance, we believe that the long-term RV stability averaged over fibers should be better than $\pm$ 2 m/s when the instrument is moved to the new vibration-controlled instrument room in fall 2007.

The Keck ET is currently being used in a pilot survey of $\sim$ 700 V = 8–12 solar-type stars in 12 different fields; 5–25 measurements have been obtained for each of the survey stars. Keck ET measurements of five known extrasolar planets are consistent with previous results (see Figure 16), and a number of possible planet/brown dwarf candidates have been identified (e.g., Figure 17).

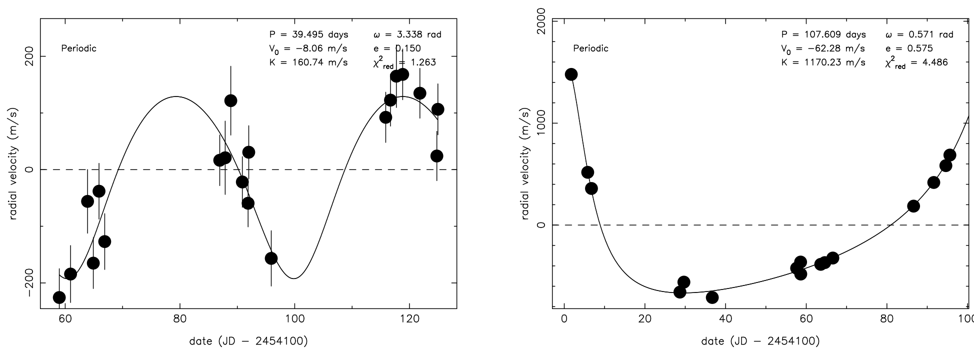

Figure 17: *Left:* An intermediate period planet candidate with $m \sin i$ of $\sim$2.5 $M_{\rm Jup}$ and $\sim$39.5 day period around a V=11.3 MS star, identified from the pilot survey data. *Right:* A rare brown dwarf candidate with $m \sin i$ of $\sim$20 $M_{\rm Jup}$ and $\sim$108 day period around a V=9.2 MS star, identified from the pilot survey data.

These early results demonstrate the ability of multi-object Doppler instruments to conduct a planet survey for short-period and intermediate-period giant planets. To achieve the main science goals of the survey, however, the instrument and software must be considerably improved. We have identified a number of hardware and software upgrades that will allow us to achieve the performance required for MARVELS.

Improvements to the instrument throughput include: (1) adding a second spectrograph to capture the second output of the interferometer, thus doubling the detected photon rate, (2) replacing the current grating with a higher efficiency VPH grating, (3) adding AR coatings to the camera optics and iodine cell, and (4) shortening the fiber length when the instrument is moved to its permanent enclosure. Further improvements to the photon-limited velocity precision will come from (5) increasing the spectral resolution to $R = 13,500$ from the current $R = 5,100$, by changing from $3''$ fibers to $2''$ fibers and employing a higher dispersion grating, (6) replacing the interferometer assembly with a better aligned assembly that reduces inhomogeneous slit illumination and image aberration. The resolution increase comes at some loss in wavelength coverage, but the net impact on Doppler precision is positive. The anticipated net impact of these modifications is a factor of 2.9 improvement in the photon-limited velocity precision.

Measurements and simulations of instrument performance suggest that inhomogeneous slit illumination and image aberration dominate the $\sim$ 11 m/s instrumental systematic errors; the new, realigned interferometer assembly should reduce these errors substantially. There are two additional sources of systematic error that require significant software upgrades: 1) The current pipeline approximates the convolution of star and iodine cell spectra as a multiplication; this has the advantage that no PSF deconvolution or instrument modeling is needed to obtain the star and iodine velocity shifts, but the approximation breaks down if the barycentric velocity correction (caused by Earth's motion around the Sun) exceeds $\sim$ 3 km/s. 2) The current pipeline does not correct for moonlight contamination. Since moonlight during bright time can contribute $\sim 0.1\%$ of the light in a $3''$ fiber centered on a $V = 12$ star, a 30 km/s offset between the Sun's velocity and the star's velocity can induce a $\sim$ 30 m/s error in the measured Doppler shift.

Our "benchmark" goals for performance of the MARVELS instruments are photon-noise errors of 21.3 m/s in a 1-hour exposure at $V = 12$, scaling as $N_{\rm photon}^{-1/2}$, and true errors (including all systematic error contributions) that are a factor of 1.5 larger than the photon-noise errors at each magnitude. Our forecasts of survey efficiency and yield in the next section are based on these benchmark assumptions, with photon-noise, systematic, and contamination contributions added in quadrature to determine the velocity error of any given observation.

6.2 Survey Strategy and Forecasts

The move from single object observations to the highly multiplexed regime allows the execution of a large, systematic census of extra-solar planets. Maximizing the scientific return of the survey involves complex trades between number of stars observed, the number of visits to each star, and the detailed cadence of the observations. To guide the survey design and forecast its discovery power, we have developed a simulation program that incorporates current observational estimates of the frequency of planetary systems, the dependence of that frequency on the chemical composition of the host star, the distribution of planet masses, orbital periods, and orbital eccentricities, and the impact of weather and lunation on the timing of the observations. The simulated velocity "measurements" are processed by automated planet-searching algorithms.

Our overarching science goal is to provide a powerful statistical sample for investigating the structure of giant planet systems. After analyzing *many* simulations of various possible MARVELS surveys, we have chosen a "baseline" survey strategy with a number of observations per target field that is expected to nearly maximize the number of planet detections around main sequence dwarf stars (given a fixed amount of telescope time). The baseline cadence observes each field about 33 times over an 18 month interval. Weather fluctuations will inevitably result in some fields being observed slightly more or less than average, but since the total number of detections as a function of the number of observations per field is near a maximum (given a fixed amount of observing time), the total number of detections is relatively insensitive to such variations.

Spreading observations of each field over the full 6-year survey duration would increase the sensitivity of MARVELS to longer period planets, but our adopted 18-month window allows us to focus on the short and intermediate period planets to which MARVELS is most sensitive, and it relaxes the requirement for very high velocity stability over long time periods. Since fields are observable at different times, the 18-month window for each field effectively breaks the full survey into three two-year blocks with different sets of fields. These three blocks allow introduction of instrument upgrades or survey strategy adjustments at two intermediate points while maintaining a well defined statistical sensitivity within each block.

For each field, about 15 of the 33 visits will be concentrated in a two month period, providing the dense sampling needed to resolve degeneracies among possible short-period orbital solutions. For the remainder of the first observing season, each field will be observed once per month. After ~ 20 observations during the first year, there may be a few possible orbital periods consistent with the observations. When the field becomes accessible again, it will be observed twice per month for six months, so as to test the candidate orbital periods, phases, and amplitudes that have been identified during the first year. This observing plan allows MARVELS to survey 108 fields. Inspection of the Tycho-2 star catalog shows 400 fields in the northern celestial hemisphere in which the number of stars brighter than $V = 12$ in the telescope's 7 deg^2 field of view exceeds the planned number of MARVELS fibers (120). Given that there are many fields suitable for MARVELS observations, we have selected a subset that have a desirable distribution on the sky and a relatively high number of stars with $V < 11$.

For our survey simulations, we have combined Tycho-2 and 2MASS observations with empirical relations to estimate the stellar properties, such as effective temperature, metallicity, probability

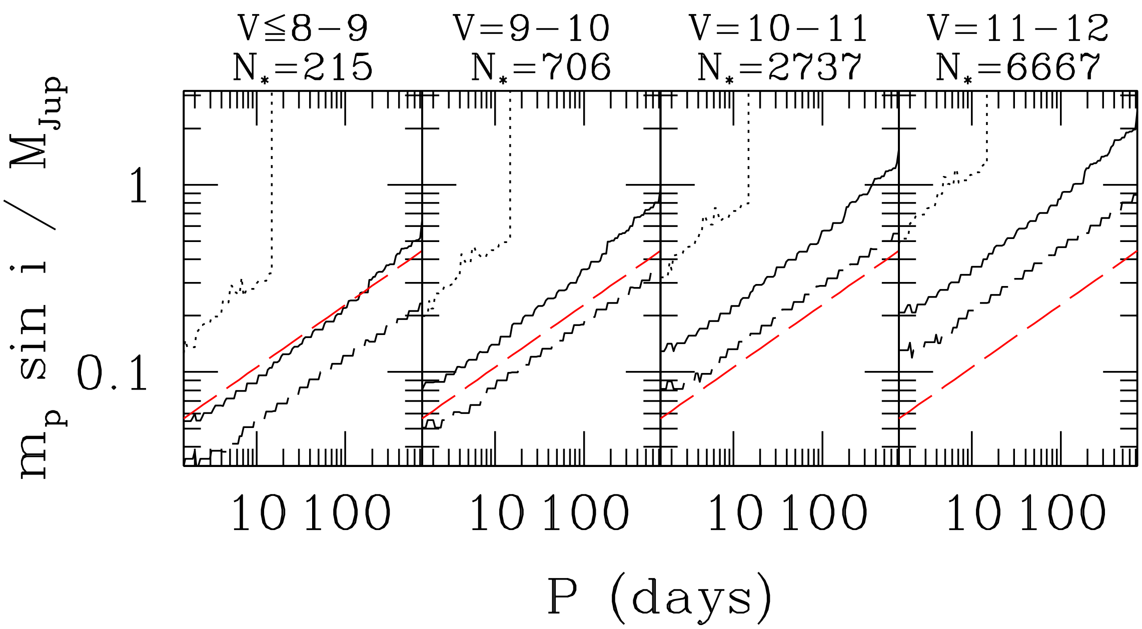

Figure 18: The efficiency of MARVELS planet detection for four different ranges of stellar apparent magnitude V. The number of stars in each magnitude range that would be monitored in the baseline survey is listed above each panel. Lines show contours of constant planet detection efficiency for host stars of the indicated brightness. At each period P, MARVELS would detect 95% of planets above the mass threshold $m_p \sin i$ indicated by the dotted line, 50% above the solid line, and 5% above the dashed line. Jitter in the curves and the jump of the dotted curve beyond 15 days arise from the finite number of stars and planets in the survey simulation. Long-dashed lines show the locus of a planet with a 10 m/s orbital velocity in a circular orbit about a solar-mass star.

of a binary companion, and probability of being a dwarf or giant (Ammons et al. 2006). Where the Ammons et al. (2006) relations result in large uncertainties, we instead adopt a distribution of stellar properties determined from a nearly volume limited sample of nearby stars. We arrive at the baseline strategy in which MARVELS will monitor approximately 11,000 stars (10,000 main sequence at $V = 8$–12 and 1000 giants at $V = 8$–11). The simulations presented below are based on real target lists, the estimated stellar properties, and simulated observing cadences (assuming random weather typical for APO).

Our Monte Carlo simulations include empirical distributions of stellar jitter as a function of stellar properties (Wright 2005), which is added in quadrature to the measurement errors listed above. We process these simulated observations with the MARVELS pipeline, which is based on a generalization of the Lomb-Scargle periodogram (Lomb 1976; Scargle 1982). A planet's orbital solution is accepted as "characterized" if (a) the false alarm probability (FAP) of the solution is less than 1×10^{-4} and (b) the FAP for the best solution is less than 0.1 times the FAP for alternative solutions with a significantly different orbital period. If the FAP is less than 1×10^{-4} but there are other viable orbital periods, the planet is "detected" but additional observations are required to provide a unique orbital solution. The FAPs are empirically calibrated from our simulations and thus include the impact of non-uniform sampling.

Figure 18 shows the efficiency of the MARVELS baseline survey at detecting planets. At

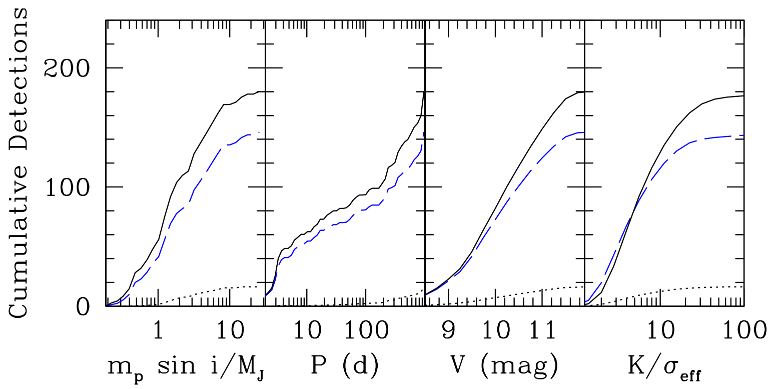

Figure 19: Forecasts of the MARVELS "planet yield" based on current estimates of planet frequency. The panels present the cumulative distributions for the expected number of planet detections as a function of planet mass, orbital period, stellar magnitude, and ratio of velocity amplitude K to single-observation rms velocity error of a comparable star with no planets. In each panel, the long-dashed curve corresponds to our benchmark assumptions about instrument and software performance and the solid curve to the case of photon-limited velocity errors. The dotted curves at bottom show, for the benchmark assumptions, planets that would require observations at other telescopes to determine the unique orbital solution (these are also included in the long-dashed curve).

$V = 10.5$, MARVELS will detect a Jupiter-mass ($M_{\rm Jup}$) planet with a period of ≈ 1 year with greater than 50% efficiency, and it will find virtually all $M_{\rm Jup}$ planets with orbital periods less than 10 days. At the faint limit of the survey, $M_{\rm Jup}$ planets must have orbital periods of less than a few months to be detected half of the time. The long-dashed lines in Figure 18 show the locus of a 20 m/s velocity amplitude (10 m/s orbital velocity) for a circular orbit, showing that unforeseen systematic error sources even at the $\sim$ 10 m/s level will affect at most a small fraction of the MARVELS sensitivity range.

Figure 19 presents the predicted yield of the MARVELS baseline survey. Dashed curves show results for the benchmark assumptions, and solid curves show the optimistic case in which velocity errors are limited by photon noise. MARVELS should find 150–200 planets around main sequence stars, with a median period of about 100 days and a median mass of $1 - 2M_{\rm Jup}$. Raising the FAP threshold to 4×10^{-3} adds ~ 80 "planet candidates" that require subsequent observations on other telescopes, about half of which would be real systems.

MARVELS will use all of the available bright time on the 2.5-m telescope until 2011, when the APOGEE program will begin observations. As discussed in §5.2, coordinated observations with shared plugplates should allow both MARVELS and APOGEE to effectively use more than 50% of the bright time thereafter. Statistical analysis of the MARVELS results requires good knowledge of the full target population, so we will obtain optical spectra of all MARVELS targets using the SDSS and (in later cycles) BOSS spectrographs.

6.3 MARVELS in Context

The extrasolar planet discoveries of the last decade have spurred a wide range of search techniques, including RV surveys, ground and space-based transit surveys, microlensing searches, astrometric

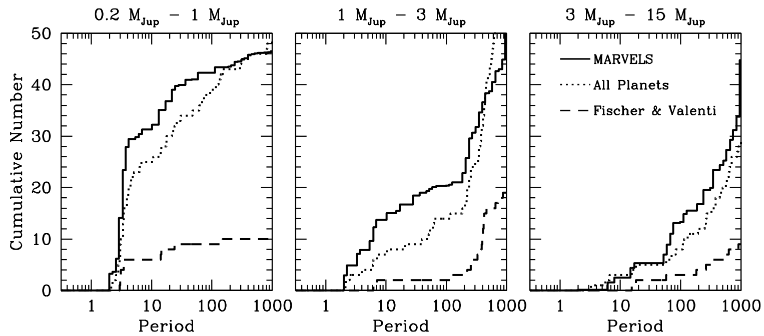

Figure 20: Cumulative number of planets as function of period in three mass ranges. The solid histograms are for planets expected from the MARVELS survey. The dotted histograms are from the current sample of nearby extrasolar planets (Butler et al. 2006). The dashed histograms are from the largest homogeneous and statistically complete sample currently available (Fischer & Valenti 2005).

searches, and direct searches for brown dwarf and planet companions. MARVELS will occupy a unique place in this landscape by providing a large sample of giant planets with short to intermediate periods. Figure 20 shows predictions for the distribution of periods of planets detected by the MARVELS survey, in comparison with the sample of 209 known nearby planets as compiled by Butler et al. (2006) and with the 46 planets from the largest statistically complete sample currently available (Fischer & Valenti 2005). Although the total planet yield from MARVELS is comparable to the known sample of planets, by focusing on massive ($M > 0.2 M_{\rm Jup}$), and short-to-intermediate period planets ($P < 1000$ days), MARVELS will significantly increase the sample of planets in this region of parameter space, which is critical for understanding the migration and dynamical evolution of giant planets. Furthermore, since these planets will be drawn from a homogeneous, well-characterized survey, the improvement in terms of true statistical power will be a factor of several over the largest comparably complete sample currently available.

The MARVELS survey will be complementary to current and planned RV surveys. The N2K (Fischer et al. 2005) and ELODIE metallicity-biased (da Silva et al. 2006) surveys are together surveying approximately 3000 stars, and they are deliberately biased toward metal-rich (and so planet-rich) targets to increase the total expected planet yield (~ 100 planets). RV surveys using instruments such as HARPS (Lovis et al. 2006) or the Rocky Planet Finder (RPF; `http://exoplanets.org/rpf.html`) generally focus their efforts in the very high precision regime, employing exquisitely crafted instruments capable of precisions and stabilities of better than 1 m s^{-1}. These surveys will principally concentrate on a moderate number of targets in order exploit the discovery space enabled by such technology (i.e., low-mass and long-period planets).

Ground-based transit surveys are primarily sensitive to short-period ($P < 5$ days), Jupiter-sized planets orbiting bright stars. However, the potential of ground-based transit surveys for studying the ensemble statistical properties of planets is limited. Extending ground-based surveys to much longer periods is impractical due to the strong scaling of the signal-to-noise ratio with period. Furthermore, interpreting the properties of a sample of planets discovered in ground-based transit surveys is extremely challenging due to the strong and difficult-to-estimate selection biases inherent in the method (see, e.g., Gaudi 2006). Spaced-based transit surveys such as *Corot* and *Kepler* will be better suited than ground-based programs to the study of ensemble planet properties, but they remain strongly biased towards short-period systems. Based on statistics from RV surveys, we estimate that *Kepler* will discover about 25 giant planets ($M \geq 0.2M_{\rm Jup}$) with periods above 5 days, with another $\sim 50-60$ at shorter periods; it may find more intermediate period, lower mass planets if these are (as expected) more plentiful than giant planets.

Astrometric and microlensing surveys are primarily sensitive to planets with longer periods than those found by MARVELS. Proposed next-generation ground and space-based microlensing surveys (Bennett et al. 2007; Gould et al. 2007) will be sensitive to Earth-mass planets with separations of several AU, but they will yield only crude orbital information. The host stars of these planets will be located at distances of several kiloparsecs, making detailed studies difficult. SIM PlanetQuest and GAIA will conduct astrometric surveys that will be complementary to MARVELS. While SIM would have the sensitivity to detect planets as small as the Earth, SIM will focus on a small number of stars ($\sim 60-240$). SIM will also observe ~ 1000 fainter "Tier II" target stars, but the sensitivity will be strongest for orbital periods approaching the mission lifetime (five years). GAIA will survey a number of stars comparable to MARVELS, but with a precision and number of observations that results in discovering a sample of planets nearly disjoint in terms of the orbital periods.

Our crucial science requirement is a large number of targets; this is needed to yield robust statistics and to provide a comprehensive probe of the dependence of planet properties on stellar properties. To carry out the planned MARVELS observations, a single-object instrument would require ~ 30 second exposures ($1/N_{\rm fiber}\times$ MARVELS exposure time). Since single-object surveys would have the option of observing brighter target stars, a high-resolution instrument like HARPS could achieve the required sensitivity even in such short exposures. However, inefficiencies due to overhead would make such a survey impossible in practice. Single-object instruments capable of high-precision and stability are best suited to intensive observations of a relatively small number of stars, in order to probe low-mass and long-period planets.

6.4 MARVELS Science

MARVELS will provide a critical data set for testing the emerging, detailed models of planet formation, migration, and dynamical evolution (e.g., Ida & Lin 2004a,b; 2005; Alibert et al. 2005; Kornet et al. 2005; Armitage 2007; Ford & Rasio 2007; Juric & Tremaine 2007; Fabrycky & Tremaine 2007; Wu et al. 2007). These models incorporate different assumptions about initial conditions and about the physical mechanisms that govern planet growth and orbital evolution. They make quantitative statistical predictions for the joint distributions of planet mass, semi-major axis, and orbital eccentricity, and for the dependence of planet frequency and orbital properties on the mass and metallicity of the host star. The short- and intermediate-period systems probed by MARVELS are especially important for understanding the physics of planet migration, one of the biggest open issues in the field of extra-solar planets. The predicted semi-major axis distribution, for example, depends on the initial radial distribution of solids and gas in the proto-planetary disk, the rate of planet migration, and the mechanisms that halt migration once it has begun. Competing models make different predictions, but the ability to test them is presently limited by the small size of well defined samples; for example, after imposing a mass completeness cut on the Fischer & Valenti (2005) sample, Armitage (2007) finds that the resulting set of 22 planets is too small to distinguish the models that he considers. MARVELS detections will be drawn from a single survey with well defined selection criteria and observing strategies. Predictive models of planet distributions can be tested quantitatively by combining them with efficiency calculations like those illustrated in Figure 18, and these tests can incorporate both detections and non-detections in a statistically rigorous fashion.

MARVELS will also test models that address the puzzling preponderance of planets with large eccentricities. Interactions with the natal gas disk are expected to damp eccentricities, so explanations of the eccentricity distribution of giant planets typically invoke dynamical processes that operate after disk dispersal. Several authors have focused on models involving planet-planet scattering (Ford & Rasio 2007; Juric & Tremaine 2007; Zhou et al. 2007), in which dynamical instabilities in multiple-planet systems lead to collisions or ejections and high eccentricities of the surviving planets. These models make specific, quantitative predictions for the distribution of eccentricities as a function of planet mass and period, which can be compared with observations.

The host stars of the ~ 150 MARVELS planet detections will be valuable targets for longer term, higher precision radial velocity monitoring to detect longer period and/or lower mass companions. This discovery plus follow-up approach is an efficient method for detecting multiple planet systems, which provide essential clues to planet formation and to the dynamical mechanisms that shape orbital distributions. For example, we would like to know what fraction of systems with migrated giant planets have surviving lower mass planets at smaller separations, or massive long-period planets. The two-stage selection effects are somewhat complicated, but they are straightforward to apply to predictive models if the follow-up program is well designed. Photometric measurements of the short-period transiting systems found by MARVELS can reveal the dynamical perturbations of additional, lower mass planets whose reflex motion may not be directly detectable. This technique is especially powerful for detecting, or limiting, the presence of companions that are captured into resonance during inward migration of the giant planets (e.g., Fogg & Nelson 2005; Zhou et al. 2005).

The broad selection of target stars will make MARVELS ideal for studying the correlation of planetary systems with stellar metallicity, mass, multiplicity, age, evolutionary stage, activity level, and rotation velocity. Because the core accretion scenario relies on the build-up of solid cores to ultimately trigger the accretion of gas in giant planets, the efficiency of planet formation is sensitive to the metallicity of the proto-planetary disk (and, by extension, the host star). Models of the core accretion scenario by Ida & Lin (2004b) predict that the planet frequency should vary roughly linearly with metallicity. An alternative scenario for the formation of gas giants, the disk instability model, predicts that the planet frequency should be nearly independent of metallicity (Boss 1997, 2002; Durisen et al. 2007). The observed correlation of planet frequency with metallicity (Santos et al. 2004; Fischer & Valenti 2005) favors the core accretion scenario, but it may be that *some* giant planets form via gravitational instability in relatively massive disks. The large number of lower metallicity stars targeted by MARVELS will help determine the relative importance of these two channels for giant planet formation.

About ten of the MARVELS planets should be transiting systems, which is more than double the number of transiting systems found to date in radial velocity surveys. This sample is small in comparison to the expected yields of ground and space-based transit surveys; however, transit surveys have strong selection effects and are biased toward bloated and short-period planets. These selection effects can be very difficult to quantify, particularly for ground-based transit surveys (Gaudi 2006). In contrast, transiting planets originally identified in RV surveys are not strongly biased towards the largest (radius) objects at fixed mass, so they are better suited to understanding the impact of host stars on planet densities.

The large number of targets makes MARVELS sensitive to rare classes of planetary systems that would be largely missed by smaller surveys: massive hot Jupiters, rapidly interacting multiple planet systems, very-hot Jupiters, and/or planets with extremely high eccentricities. MARVELS will be the best survey to date for exploring the "brown dwarf desert," the apparent paucity of $\sim 15 - 80 M_{\rm Jup}$ companions to solar-type stars. This gap may mark the transition between the most massive objects formed in circumstellar disks and in fragmenting molecular cloud cores, or it could be an indication of a mass-dependent orbital migration mechanism. If the history of planet searches is any guide, then the most interesting discoveries from MARVELS will be complete surprises, planetary systems with properties that have not previously been seen and have not been anticipated by theory.

7 Central Infrastructure

There is a well-developed, efficient infrastructure at the observatory for the support of the SDSS surveys. The SDSS-III program, with four instruments, is more complex, and we face new challenges in the areas of plate plugging and fiber mapping. We intend to keep the present system with as few changes as possible, but some additional capabilities are required.

First, the system needs to handle more cartridges. We currently have nine, each with 640 SDSS fibers and 64 MARVELS pilot-program fibers. The cartridges cannot handle more than about 1000 fibers, and since BOSS uses 1000 fibers, we will build new cartridges so we have enough to service SEGUE-2, MARVELS, and APOGEE. We will implement a new (and simpler) cartridge handling system.

The BOSS cartridges are much like the current SDSS ones, with short fibers leading from ferrules plugged into the plates to fixed fiber slits that latch to the spectrographs; the only significant change is the number of fibers, from 640 to 1000. The fibers are mapped (determining which fiber goes to which hole) by a machine that can be trivially adapted to the BOSS cartridges and to the SEGUE-2 fibers in the other cartridges.

MARVELS and APOGEE, on the other hand, both use fixed instruments a long way from the telescope and so will both use long fixed fiber runs. These will be mated to the interchangeable cartridges by sets of fiber couplers, each handling 32 fibers, which will be arranged in one robust, blind-mating connector for each instrument. The losses in these couplers can be kept to about 10%. The current mapping scheme cannot be used as is for these fibers, since there is no fiber slit. However, it is a simple matter to make a dummy slit with identical connectors, which will allow mapping with the same machine and software used for BOSS and SEGUE-2.

In the current SDSS surveys it does not matter which fiber goes to which hole — the mapping sorts out the identifications, and that is all that matters. This is still true for APOGEE, but for the other surveys some fiber management will be necessary, varying in severity from BOSS, for which quasar spectra should not be adjacent to each other or to bright standards, to bright-time SEGUE-2, for which every other fiber is sky, to MARVELS, where achieving the desired accuracy demands that the same fiber go to the same star in repeat visits.

Since the fibers are plugged by hand by two individuals working at the same time, a reasonable requirement is that the plugging and mapping of a full night's complement of cartridges can be done during a working day. This requires a system of fiber management which is simple and robust. The scheme we currently plan to implement is based on color-coding the fibers and holes in groups of ten. Both fibers and (for MARVELS) fiber groups will be color coded, and each group will be physically bound for easy identification and separation from its neighbors. Color-coding the fibers is easy. To color-code the *holes*, we plan to use a large flatbed plotter with which the plugging patterns and codes are marked directly on the aluminum plates before use. We believe that this will be satisfactory, but we are also exploring alternative methods.

8 Data processing

8.1 BOSS, SEGUE-2 and APOGEE data reductions

The analysis of the BOSS imaging data will be done at Princeton using the SDSS `photo-op` pipeline, which has been run on all SDSS imaging data to date. In addition to running the photometric reductions (Lupton et al. 2001), this pipeline resolves multiple detections of objects, determines the window function of the imaging survey, and performs the critical "ubercalibration" step to produce a homogeneous zeropoint across the entire survey (Padmanabhan et al. 2007). The results of these photometric reductions have formed the basis of most of the large-scale structure results from SDSS to date (e.g., Tegmark et al. 2004; Eisenstein et al. 2005; Zehavi et al. 2005; Tegmark et al. 2006). Final target selection for BOSS will proceed from the integrated set of SDSS-I, II, and BOSS imaging that will be available early in 2009.

Data reduction from the BOSS and APOGEE spectrographs consists of two largely independent steps: extraction and classification. The extraction step from the raw CCD images results in wavelength-calibrated and flux-calibrated spectra. The BOSS and SEGUE-2 data reduction will build directly upon SDSS pipelines `idlspec2d` and `idlutils`. These are flexible and modular, employ sophisticated algorithms where necessary, and have proven quite successful for SDSS. Due to the transparency and good documentation of these codes, they have comprised the critical elements of analysis pipelines of other surveys, including the Deep Extragalactic Evolutionary Probe (DEEP) at Keck and Hectospec at the MMT.

The extraction step requires more work for each of the surveys. BOSS will require handling fainter targets relative to SDSS. It will need to achieve at least the current 2% error in sky-subtraction, to improve upon the current 4% accuracy in modeling the noise in the QSO Lyα forest spectra, and to mitigate sky-noise coupling between the QSO fibers. These improvements will require the development of extraction algorithms that use 2-dimensional PSFs and that are much more analogous to image analysis codes. However, to facilitate commissioning of the instrument and target selection early in the survey, we will begin by minimally modifying the existing pipelines to handle the new data format from the BOSS spectrographs. As we develop algorithmic improvements, we will incorporate them as module replacements to the existing software. The bright time operations of SEGUE-2 after installation of the BOSS spectrograph, described in §4, will require handling of the brighter sky background. Finally, the APOGEE extraction needs to be created: it will draw on these tools and on the SDSS experience, but it will need to be optimized for the near-infrared, very high resolution regime.

The classification step fits these extracted, 1-D spectra to template spectra, and derives other quantities such as redshifts and velocity dispersions for galaxies, and metallicities, abundances, surface gravities, and temperatures for stars. The final product consists of the 1-D spectra and their derived parameters. Again, the classification step will draw heavily on the tools SDSS and SEGUE developed. Changes will be required for BOSS and the SEGUE-2 bright time data, to adjust to the new instrument's resolution and signal-to-noise properties (in order to assure consistent results with the SDSS spectrographs). For the APOGEE classification step, no comparably ambitious project has ever been undertaken, and we anticipate significant effort in producing well calibrated, well understood parameter estimates. Our procedure will likely follow that used recently by Rich & Origlia (2005) and Cunha & Smith (2006), iteratively fitting the parameters to the results of a combination of stellar evolution and stellar atmosphere codes. Although these methods are a sensible starting point, such analyses are the subject of cutting-edge research and the methods used are still being refined. The data releases will of course include the underlying spectra, since we expect that astronomers outside the collaboration will hone their own techniques on this massive data set (as they have with previous SDSS data).

8.2 MARVELS data reductions

The MARVELS data processing will build on the pipeline already created for the ET pilot project, which extracts the phase shift from the fringe images for each fiber, combines observations of the same objects over many nights to produce a radial velocity curve as a function of time, and searches that curve for a periodic signal (see Figure 16). To reach the survey radial velocity precision goal of a few m s^{-1}, we are modifying the code to optimally extract the signal from the fringe images produced by the instrument, to properly treat data at the fringe edges, to model image distortions, to fully model the convolution of stellar and iodine spectra, to mitigate the effects of moonlight contamination (particularly important for the faintest stars in the survey), and to disentangle the (low level of) cross-talk between fibers. These changes will significantly increase the S/N in the processed spectrum as well as reducing the systematic errors in the measurements.

The science of radial velocity curve analysis is already quite advanced, and there are a number of high-quality software packages in use. MARVELS team members have already created planetary search code capable of handling MARVELS observations, and this will continue to undergo refinement as the survey progresses. The public releases will include the periodograms for each of the stars, but for those researchers (probably the majority) who wish to use their own software to analyze the radial velocity measurements, we will provide the radial velocity curves as well. We will also provide spectra of the target stars from the SDSS or BOSS spectrographs, and derived stellar parameters.

8.3 Data integration

The data from the four survey pipelines will be stored in the Science Archive Servers, in a central location at NYU, in a documented format, of the sort that this team developed in the past for SDSS-I and II. For BOSS imaging and BOSS spectroscopy, we will link the imaging, targeting, plate design, and spectroscopic results into an integrated package that appropriately matches the imaging catalog to the spectroscopic catalogs and that also provides a fully-specified window function and completeness map of the survey. Because the scientific goals of measuring large-scale structure will rely on access to already-determined redshifts from SDSS-I and II, we will incorporate those results into a master redshift catalog. Our model for integrating these data sets and for building the window function is the NYU Value-Added Galaxy Catalog (Blanton et al. 2005), which served this function for the SDSS and has been heavily used both inside and outside the SDSS Collaboration. SEGUE-2 and APOGEE targeting and spectroscopy will be cross-identified with the BOSS imaging where appropriate. They will also be cross-identified with each other in order to allow easy comparison of the stellar parameter estimates between the two surveys.

8.4 Public access to data

The high-level data processing schedule is summarized briefly in Table 2, including the timing of the survey data releases. We plan for five public data releases, in the summers of 2010 to 2013, with a final release at the end of 2014. These are denoted as DR9 to DR13, in keeping with the SDSS and SDSS-II nomenclature. Because of the long-duration nature of the MARVELS experiment, there will be just three MARVELS data releases, tied to DR10, DR12 and DR13. For other surveys, the data releases will be similar in form to the current SDSS releases. We will distribute all data, including BOSS imaging data, using the powerful Catalog Archive Server and SkyServer interfaces developed for SDSS-I and II, as well as through the Science Archive Servers to provide access to the images and "flat file" versions of the catalogs. Our team includes most of the core group that developed the SDSS versions of these tools, which are extremely popular and have helped make the SDSS such a high-impact survey in astronomy and cosmology.

The Catalog Archive Server (CAS) and the accompanying CASJobs query server provide a powerful tool for exploring the data and in particular for finding rare objects quickly. The servers provide an easy-to-use front web page, and they allow more expert users to write complex SQL queries. The CAS is compatible with Virtual Observatory (VO) tools — it is an OpenSkyNode, and it is in fact designed by many of the principals in the VO effort. The number of distinct, valid queries that have been launched in the SDSS database in the past five years or so is staggering, around 7 million; these queries have returned a total of about 100 billion rows of information.

The CASJobs interface is a particularly powerful component of the CAS, especially compared to other comparable efforts in astronomy. It gives users a staging area called "MyDB" to manipulate and refine their query results before sending the results back over the network, saving considerable time and energy for user. Furthermore, it allows users to import their own tables into MyDB, join them with tables in the SDSS database, and to share data within user groups. We can identify over 500 individual, "heavy-duty" astronomer-users of the CASJobs database, demonstrating the popularity and usefulness of this research tool. For all these reasons, we plan to continue using CAS and CASJobs as our primary data distribution mechanism.

The SkyServer has been highly successful in providing data to both expert and novice users. It provides a graphical interface to the data set that allows navigation across the sky, a natural interface for bringing up ancillary data on objects in the imaging, such as spectra and matches to external databases, and the easy production of finding charts. As we outline in the EPO section, because of its ease of use, the SkyServer is useful both as a research tool and as an interface to the data for the public. Although the SkyServer is designed to be relatively independent of the underlying database, the changes in CAS outlined above will require some software effort to make SkyServer compatible.

We have budgeted for three copies of this database to maximize the speed of access across the globe: one master copy in the United States, one copy in Europe, and one copy in Asia. For each survey, we plan to also release the integrated "flat files" in the Science Archive Servers described in §8.3. Experience in the SDSS has shown that access to the catalogs in this form can be more convenient for certain science than the databases, and providing access to these files is low in effort. These servers also allow users to access the digital images that the catalogs and spectra are based upon, which in our experience many "high-end" users are interested in re-measuring. The Science Archive Servers will be run similarly to the SDSS Data Archive Servers and the NYU Value-Added Galaxy Catalog. They will be published in VO and allow cone searches as well as Simple Imaging Access Protocol (SIAP) queries for imaging data.

9 Education and Public Outreach

The unique and far reaching education and public outreach (EPO) programs of SDSS-I and SDSS-II have been a great success, supporting a broad range of science education for diverse audiences. We will extend this successful model to SDSS-III and expand our efforts into new educational activities.

Projected date	Data release	APOGEE	BOSS	MARVELS	SEGUE-2
July 2009	—	—	Photometric catalog (for targeting)	—	—
July 2010	**DR9**	—	**Photometric catalog (complete)**	—	**Dark-time spectra (complete)**
July 2011	**DR10**	—	**Spectroscopic catalog (up to July 2010)**	**Radial velocity curves (up to July 2010)**	**Bright-time spectra (up to July 2010)**
July 2012	**DR11**	—	**Spectroscopic catalog (up to July 2011)**	—	Bright-time spectra (up to July 2011)
July 2013	**DR12**	**Spectroscopic catalog (up to July 2012)**	**Spectroscopic catalog (up to July 2012)**	**Radial velocity curves (up to July 2012)**	**Bright-time spectra (up to July 2012)**
July 2014	—	Spectroscopic catalog (up to July 2013)	Spectroscopic catalog (up to July 2013)	—	Bright-time spectra (up to July 2013)
Dec. 2014	**DR13**	**Spectroscopic catalog (complete)**	**Spectroscopic catalog (complete)**	**Radial velocity curves (complete)**	**Bright-time spectra (complete)**

Table 2: Data processing milestones for SDSS-III, as a function of time for each of the four surveys. Bold entries indicate milestones that will be included in that year's data release. For data releases other than DR13, we expect to have the relevant catalogs reduced nine months before release and loaded into the CAS databases six months before release.

We will leverage existing and emerging educational technologies, including video-podcasting, PDA-based interfaces, and immersive virtual environments like SecondLife (`www.secondlife.com`), to develop activities that will meet our goal of bringing usable, professional quality data to support formal and informal education.

One of the features that sets the SDSS EPO efforts apart from similar projects is that we have always made our data easily accessible to a wide variety of learners — anyone with a web connection can access the same data that the professional researchers use. In 2001, early in the data distribution phase of the SDSS-I project, we created the SkyServer web site (`http://skyserver.sdss.org`), which provides public-friendly tools to browse and view SDSS data. Programmers at the Johns Hopkins University (JHU) created web interfaces to display and search the data. Jordan Raddick, a science education and public outreach coordinator at JHU, worked with a master high school teacher to create and field-test inquiry-oriented classroom lesson plans on SkyServer. For example, one project asks students to create a Hertzsprung-Russell diagram by searching for and plotting SDSS magnitude data for star clusters. Already more than 100 teachers, primarily at the high school and "Astronomy 101" undergraduate level, have been trained to use the SDSS tools and projects.

SDSS-I and SDSS-II have also achieved success in the realm of informal education. For example, the American Museum of Natural History (AMNH) devoted one of its *Science Bulletins* entirely to the SDSS (`http://sciencebulletins.amnh.org/astro/f/sdss.20051208/index.php`), and there are SDSS-based exhibits at AMNH, the National Air and Space Museum, Adler Planetarium, and other major science museums. Several museums have made extensive use of SDSS data as part of their regular planetarium shows.

SDSS-II has a single EPO officer, whose primary role is to coordinate the outreach efforts that other teams develop with their own funding. This approach has allowed SDSS-II to leverage and amplify its outreach efforts by supplying data and expertise into proven EPO programs with existing networks of users. We will continue this successful model as we develop an EPO program for SDSS-III, which will offer an infusion of new data, together with a team of EPO and technical developers who know how to turn the data into a form that learners can use. SDSS-III can offer activities to address many topics in a dynamic astronomy curriculum, including spectroscopy and chemical composition, sky observations, extra-solar planets, the structure of the Milky Way galaxy, quasars and black holes, the formation of cosmic structure, and the matter and energy contents of the Universe.

In SDSS-I and SDSS-II, data display and search tools were a central feature of our EPO programs. Building on this base, we will work with a programmer at JHU to develop new access tools, adapting from existing SkyServer tools where possible. In particular, we will create new tools to more intuitively display spectra, and to work with repeat observations of the same objects. We will explore distributing data using the emerging technologies mentioned above. Educators have already used these technologies effectively in courses, and we expect that distributing data in these new ways will lead to new educational possibilities.

As with SDSS-II, one of the major duties of the SDSS-III EPO officer will be to seek out partnerships with other teams of outreach developers. The first of our EPO partners will be members of the Center for Astronomy Education (CAE) and the Conceptual Astronomy and Physics Education Research (CAPER) team at the University of Arizona. Individuals working in CAE and CAPER are well known in the astronomy education community for their development of unique and successful instructional materials and strategies, which have been shown through research to be effective in the teaching and learning of astronomy. For the current project, postdocs and graduate students with expertise in astronomy education, working under the supervision of faculty members from CAE and CAPER, will help the EPO team to develop curriculum materials for K-12 and college classrooms aligned with the National Research Council's National Science Education Standards, and to conduct field-testing in both K-12 and introductory undergraduate environments using their extensive networks of educators. Julie Lutz of University of Washington, the EPO Director for SDSS-II and a part of CAE and CAPER's northwest regional teaching exchange, will also work with SDSS-III as an education consultant and developer. She will consult with SDSS-III primarily on strategies for K-12 formal education.

We will spend our development resources primarily on formal education, where we have had so much success in the past. We propose a two-pronged approach to our formal education activities: long-term, inquiry-based activities to be integrated into a curriculum, and shorter, focused activities on a single topic, to be used as single labs or homework assignments. To meet the need that so many teachers have expressed for authentic student experiences, every activity that we produce will be designed to engage students and teachers in investigations with real SDSS data.

The audience for the long-term activities will be instructors and students in inquiry-based astronomy or integrated science courses. Some settings for these courses include college-level astronomy or integrated science for non-majors, science methods for preservice teachers, and mathematics or technology courses. The primary instructional design for developed lessons will be in the form of authentic case studies. Students will be presented with rich contextual problems framed as open-ended scientific research questions, and will be facilitated to investigate with our tools to develop their own evidence-based conclusions. Some possible topics for case studies include the power source of quasars, the history of Milky Way assembly, the properties of extra-solar planetary systems, and the expansion of the Universe. Students will use SDSS data to formulate an evidence-based argument about their scenario presented in their case study. These case-study analyses are designed to take $3-4$ weeks for students to complete. We will emphasize in our teacher materials that the most important part of assessing the activity is not whether students get the "correct" (i.e., currently accepted by astronomers) answer, but the degree to which they defend their conclusions with evidence-based scientific reasoning using authentic SDSS data.

Although the case study approach we describe is shown by educational research to be particularly effective in teaching science, we recognize that some instructors will not have enough time to devote to such an approach. To meet the needs of these instructors, we will develop shorter activities that can be completed in a single $2-3$ hour lab period or can be done as individual student homework. These activities will use guided inquiry as much as possible. Potential topics include the H-R diagram, the motions of stars in the Milky Way, and analyzing stellar radial velocities to discover extra-solar planets.

Although we plan to focus primarily on formal education, we will also develop an EPO program for informal education. We expect to continue our role supplying data and development expertise to planetariums, especially those we have worked with in the past. We will also plan curricula for students in summer science camp settings, developing a model that can be readily implemented in many locations in the form of a "science camp in a box." These activities will be adapted wherever possible from our formal education activities, including developing extensions and further research ideas for motivated science camp attendees.

Prior to the first SDSS-III public data release in 2010, we will develop new tools to intuitively display spectra and repeat observations, and we will write background explanations of the science and technology behind SDSS-III. During this phase, we will continue to offer training sessions for teachers with educational materials from SDSS-I and SDSS-II, so that when new lessons become available we will have a solid network of potential adopters. We will develop one pilot project from each formal education approach, so that we can have "first light" of EPO materials on the SDSS-III website on the day of the first SDSS-III public data release. We will continue to develop new educational activities and will continue to offer training in the projects we create.

10 Collaboration and Personnel

10.1 Membership

The SDSS-III collaboration will be modeled after the successful structure of SDSS-I and SDSS-II. Institutions will join the collaboration via contributions, both technical and financial. Scientists at these institutions will have data rights to all of the SDSS-III surveys. "Full membership" will yield data rights for all scientists at an institution. "Associate membership," which will require a smaller contribution, will yield data rights for only a limited number of scientists. More than 20 institutions have formally expressed interest in joining SDSS-III, and we continue to actively seek new institutional members.

The SDSS-III will have an Advisory Council to oversee the Director and the project. The Council will be comprised of representatives of the member institutions and will itself report to the ARC Board of Governors. The Advisory Council will be formed after the first round of institutions sign MOUs. Until that time, this role is filled by a Steering Committee.

SDSS-III is managed by the Central Project Office, chaired by the Director and including the Program Manager, Project Scientist, Technical Coordinator, Data Coordinator and Survey Coordinator. Each of the individual surveys has its own PI and associated management team, including survey and instrument scientists. The PIs report to the Central Project Office, which has final budget, schedule and technical authority in all matters. The organizational chart is shown in Fig. 21. Nearly all of the key positions have been filled, and substantial teams of scientists are already at work developing the four surveys.

10.2 Personnel

The program described here emerged out of a solicitation by the Astrophysical Research Corporation of proposals for use of the APO 2.5-meter telescope and its instruments after the completion of SDSS-II. Each survey has a substantial team of scientists who have been working over the past year on instrument and survey design, including many investigators who have been involved since the original proposals.

The "high-level organizational chart" (Fig. 21) includes the SDSS-III Management Committee and Steering Committee and, in addition, the Survey Scientists and Instrument Scientists, the APO Site Operations Manager, the Lead Observer, and the EPO Coordinator (see Fig. 21). These individuals are: Daniel Eisenstein (Director), Timothy Beers (SEGUE-2 Survey Scientist; Steering Committee Member), Michael Blanton (Data Coordinator), Holland Ford (MARVELS Survey Scientist; Steering Committee Member), Jian Ge (MARVELS PI), Bruce Gillespie (Program Manager), James Gunn (Technical Coordinator), Mark Klaene (APO Site Operations Manager), Richard Kron (Steering Committee Chair), Steven Majewski (APOGEE PI and Acting Survey Scientist), Robert O'Connell (Steering Committee Member), Jordan Raddick (EPO Coordinator), Constance Rockosi (SEGUE-2 PI and Instrument Scientist), Natalie Roe (BOSS Instrument Scientist; Steering Committee Member), David Schlegel (BOSS PI), Michael Skrutskie (APOGEE Instrument Scientist),

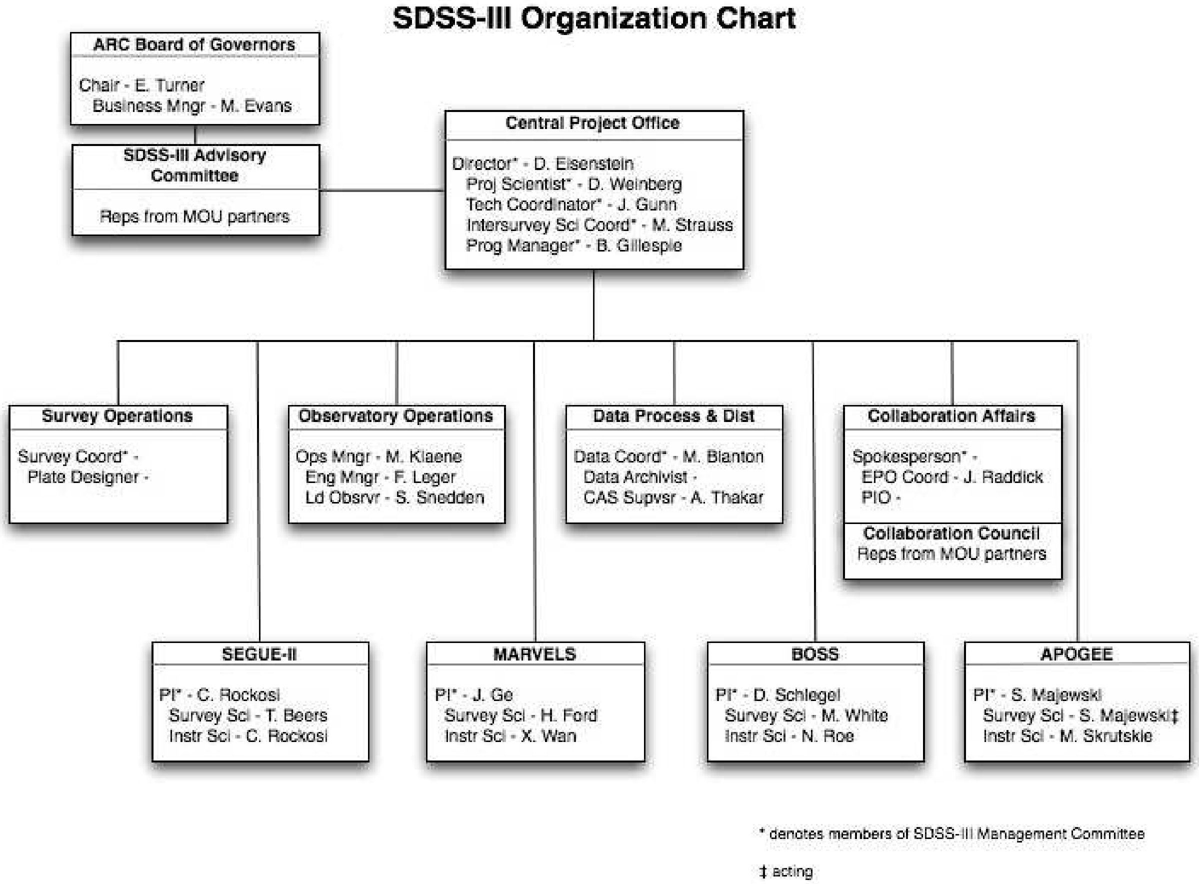

Figure 21: Top-level organizational chart of SDSS-III.

Stephanie Snedden (Lead Observer), Michael Strauss (Intersurvey Science Coordinator), Xiaoke Wan (MARVELS Instrument Scientist), David Weinberg (Project Scientist), Martin White (BOSS Survey Scientist).

Each survey team has appointed a "Project Editor," who organizes contributions to the SDSS-III proposals and documentation (including this project description). The four Project Editors are Jill Knapp (SEGUE-2), Ricardo Schiavon (APOGEE), Donald Schneider (MARVELS), and Martin White (BOSS). In addition to the editors and those listed above, others who contributed calculations or text in direct support of this document include: Eric Agol, Robert Cahn, Eric Ford, Scott Gaudi, Andrew Gould, Jennifer Johnson, Stephen Kane, Guinevere Kauffmann, Juna Kollmeier, Suvrath Mahadevan, Patrick McDonald, Heather Morrison, Nikhil Padmanabhan, Gordon Richards, Sara Seager, Uros Seljak, and John Wilson. We have also benefited from review and comments by many other members of the SDSS-III Collaboration and by a number of external reviewers who generously contributed their time and expertise.

References

Abadi, M.C., Navarro, J.F., Steinmetz, M., & Eke, V. 2003a, ApJ, 591, 499

Abadi, M.C., Navarro, J.F., Steinmetz, M., & Eke, V. 2003b, ApJ, 597, 21

Abadi, M.C., Navarro, J.F., & Steinmetz, M. 2006, MNRAS, 365, 747

Abbott T., et al., The Dark Energy Survey, white paper submitted to the Dark Energy Task Force [astro-ph/0510346]

Albrecht A., et al, Report of the Dark Energy Task Force, 2006 [astro-ph/0609591]

Albrecht A., Skordis C., 2000, Phys. Rev. Lett., 84, 2076.

Alcock, C., & Paczyński, B. 1979, Nature, 281, 358

Alibert, Y., Mordasini, C., Benz, W., & Winisdoerffer, C. 2005, A&A, 434, 343

Allende Prieto, C., et al. 2006, ApJ, 636, 804

Ammons, S. M., Robinson, S. E., Strader, J., Laughlin, G., Fischer, D., & Wolf, A. 2006, ApJ, 638, 1004

Angulo R., Baugh C.M., Frenk C.S., Lacey C.G., 2007, preprint, [astro-ph/0702543]

Arkani-Hamed N., Dimopoulos S., Dvali G., 1998, Phys. Rev. Lett., B429, 257

Armendariz-Picon C., Mukhanov V., Steinhardt P.J., 2000, Phys. Rev. Lett. 85, 4438

Armitage, P.J. ApJ in press/astro/ph-arXiv:0705.3039

Battaglia, G., et al. 2005, MNRAS, 364, 433

Bebek C.J., et al., 2004, SPIE 5499, 10

Becklin, E.E., et al. 2005, ApJ, 632, L119

Beers, T.C., Rossi, S., Norris, J.E., Ryan, S.G., & Shefler, T. 1999, AJ, 117, 981

Bell, E.F., et al. 2007, ApJ (in press: arXiv:0706.0004)

Belokurov, V., et al. 2006a, ApJ, 642, L137

Belokurov, V., et al. 2006b, ApJ, 647, L111

Belokurov, V., et al. 2007, ApJ, 654, 897

Bennett, D. P., et al. 2007, ExoPTF White Paper (arXiv:0704.0454)

Bensby, T., et al. 2005, A&A, 433, 185

Blake C., Glazebrook K., 2003, ApJ, 594, 665

Blake C., Collister A., Bridle S., Lahav O., 2007, MNRAS, 374, 1527

Blanton, M. R. et al. 2005, AJ, 129, 2562

Bolton, A. S., Burles, S., Koopmans, L. V. E., Treu, T., & Moustakas, L. A. 2006, ApJ, 638, 703

Boss, A. P. 1997, Science, 276, 1836

Boss, A. P. 2002, ApJ, 567, L149

Bousso R., Rev. Mod. Phys., 2002, 74, 825

Brewer, M., & Carney, B.W. 2006, AJ, 131, 431

Brook, C.B., Kawata, D., Gibson, B.K., & Freeman, K.C. 2004, ApJ, 512, 894

Brook, C.B., Gibson, B.K., Martel, H., & Kawata, D. 2005, ApJ, 630, 298

Brown, W.R., Geller, M.J., Kenyon, S.J., & Kurtz, M.J. 2005, ApJ, 647, 303

Brown, W.R., Geller, M.J., Kenyon, S.J., Kurtz, M.J., & Bromley, B.C. 2007, ApJ, 660, 311

Bullock, J.S., & Johnston, K.V. 2005, ApJ, 635, 931

Bullock, J.S., Kravtsov, A.V., & Weinberg, D.H. 2001, ApJ, 548, 33

Butler, R.P. et al. 1996, PASP, 108, 500

Butler, R.P., et al. 2004, ApJ, 617, 580

Butler, R.P., et al. 2006, ApJ, 646, 505

Cannon R., et al., 2006, MNRAS, 372, 425

Capozziello S., 2002, Int. J. Mod. Phys., D11, 483

Carroll S.M., Duvvuri V., Trodden M., Turner M.S., 2004, Phys. Rev. D70 043528

Cen R., Miralda-Escudé J., Ostriker J.P., Rauch M., 1994, ApJ, 437, L9

Charbonneau, D., Brown, T.M., Burrows, A., Laughlin, L. 2007, "When Extrasolar Planets Transit Their Parent Stars", in Protostars and Planets V, B. Reipurth, D. Jewitt, & K. Keil, eds. (Tucson: University of Arizona Press)

Christlieb, N., et al. 2002, Nature, 419, 904

Clewley, L., Warren, S.J., Hewett, P.C., Norris, J.E., & Evans, N.W. 2004, MNRAS, 352, 285

Clewley, L., et al. 2005, MNRAS, 362, 349

Coble K., Dodelson S., Frieman J.A., 1997, Phys. Rev. D55, 1851

Crane, J.D., et al. 2003, ApJ, 594, L119

Crocce M., Scoccimarro R., 2007, preprint [arvix/0704.2783]

Croft R.A.C., et al., 1998, ApJ, 495, 44

Cunha, K. & Smith, V. V. 2006, ApJ, 651, 491

Cushing, M. C. Rayner, J.T., Davis, S.P. & Vacca, W.D. 2003, ApJ, 582, 1066 da Silva, R., et al. 2006, A&A, 446, 717

Dehnen, W. 1998, AJ, 115, 2384

Dehnen, W., McLaughlin, D., & Sachania, J. 2006, MNRAS, 369, 1688

De Silva, G.M., Freeman, K.C., Bland-Hawthorn, J., Asplund, M., & Bessell, M.S. 2007, AJ, 133, 694

Diemand, J., Kuhlen, M., & Madau, P. 2007, ApJ (in press: astro-ph/0703337)

Dufour, P., et al. 2007, ApJ, 663, 1291

Durisen, R.H., Boss, A.P., Mayer, L., Nelson, A., Rice, K., & Quinn, T.R. 2007, "Gravitational Instabilities in Gaseous Protoplanetary Disks and Implications for Giant Planet Formation", in Protostars and Planets V, B. Reipurth, D. Jewitt, & K. Keil, eds. (Tucson: University of Arizona Press)

Einstein A., 1917, Sitzungsberichte der Preussischen Akad. d. Wissenschaften, 142

Eisenstein D.J., Hu W., Tegmark M., 1998, ApJ, 504, L57

Eisenstein D.J., 2003, in *Wide-field Multi-Object Spectroscopy*, ASP Conference Series, ed. A. Dey

Eisenstein, D. J., Blanton, M., Zehavi, I., Bahcall, N., Brinkmann, J., Loveday, J., Meiksin, A., & Schneider, D. 2005, ApJ, 619, 178

Eisenstein D.J., et al., 2005, ApJ, 633, 560

Eisenstein D.J., 2005, New Astronomy Reviews, 49, 360.

Eisenstein D.J., Seo H.-J., White M., 2007, ApJ, 664, 660

Eisenstein D.J., Seo H.-J., Sirko E., Spergel D.N., 2007a, ApJ, 664, 675

Eisenstein, D.J., Dey, A., Glazebrook K., Nichol, B., Parkinson, D., White, M., "Report on the Dark Energy Science Case for WFMOS", submitted to the Gemini Board, May 3, 2007.

Eggen, O.J, 1996, AJ, 112, 1595

Erskine D., & Ge, J. 2000, in Proc. Imaging the Universe in Three Dimension, Edited by W. van Breugel and J. Bland-Hawthorn, ASP Conference Series, 195, 501

Erskine, D., 2003, PASP, 115, 255

Fabrycky, D., & Tremaine, S. 2007, ApJ, submitted (arXiv:0705.4285)

Fellhauer, M., et al. 2006, ApJ, 651, 167

Fischer, D.A. & Valenti, J. 2005, ApJ, 622, 1102

Fischer, D.A., et al. 2005, ApJ, 620, 481

Flynn, C., Holmberg, J., Portinari, L., Fuchs, B., & Jahreiß, H. 2006, MNRAS, 372, 1149

Fogg, M. J., & Nelson, R. P. 2005, A&A, 441, 791

Font, A.S., Johnston, K.V., Bullock, J.S., & Robertson, B.E. 2006a, ApJ, 638, 585

Font, A.S., Johnston, K.V., Bullock, J.S., & Robertson, B.E. 2006b, ApJ, 646, 886

Ford, E. B., & Rasio, F. A. 2006, ApJ, 638, L45

Ford, E. B., & Rasio, F. A. 2007, ApJ, submitted (astro-ph/0703163)

Frebel, A., et al. 2005, Nature, 434, 871

Freeman, K., & Bland-Hawthorne, J. 2002, ARAA, 40, 487

Fuhrmann, K. 1998, A&A, 338, 161

Fukugita, M., Hogan, C. J., & Peebles, P. J. E. 1998, ApJ, 503, 518

Fulbright, J. P., McWilliam, A., & Rich, R. M. 2006, ApJ, 636, 821

Gänsicke, B.T., Marsh, T.R., Southworth, J., & Rebassa-Mansergas, A. 2006, Science, 314, 1908

Gänsicke, B.T., Marsh, T.R., & Southworth, J. 2007, MNRAS (in press: arXiv:0705.0447)

Gates, E., et al. 2004, ApJ, 612, L129

Gaudi, B. S. 2006, ASP Conf. Series: "Transiting Extrasolar Planets Workshop" (arXiv:astro-ph/0612141)

Ge, J., 2002, ApJ, 571, L165

Ge, J. et al. 2002, PASP, 114, 1016

Ge, J., et al. 2006a, ApJ, 648, 683

Ge, J., et al. 2006b, Proc.SPIE, 6269, 75

Gerhard, O. 2002, Space Science Reviews, 100, 129

Gilmore, G., Wyse, R.F.G., & Norris, J.E. 2002, ApJ, 574, L39

Glazebrook, K., et al., 2007, "The WiggleZ project: AAOmega and dark energy", Proceedings of the Durham "Cosmic Frontiers" ASP conference eds. Metcalfe & Shanks [astro-ph/0701876]

Gnedin, O.Y., Gould, A., Miralda-Escudé, J., & Zentner, A.R. 2005, ApJ, 634, 344

Gould, A., Dorsher, S., Gaudi, B. S., & Udalski, A. 2006, Acta Astronomica, 56, 1

Gould, A., Gaudi, B. S., & Bennett, D. P. 2007, ExoPTF White Paper (arXiv:0704.0767)

Gould, A., & Morgan, C. W. 2003, ApJ, 585, 1056

Governato, F., et al. 2007, MNRAS, 374, 1479

Harris, H,C., et al. 2006, AJ, 131, 571

Harris, H.C., et al. 2007, submitted to AJ

Hartwick, F.D.A. 1987, in "The Galaxy", D. Reidel Publishing Co., p. 281

Hatzes, A.P., & Cochran, W.D., 1992, ESO Workshop on High Resolution Spectroscopy with the VLT. Proceedings, held in Garching, Germany, February 11-13, 1992. Editor, M.-H. Ulrich; Publisher, European Southern Observatory, Garching bei Munchen, Germany, 1992. P. 275

Helmi, A. 2004, ApJ, 610, L97

Helmi, A., et al. 2003, ApJ, 586, 195

Helmi, A., & White, S.D.M. 1999, MNRAS, 307, 495

Helmi, A., et al. 2006, MNRAS, 365, 1309

Hill G., et al., 2004, Mitchell Symposium: "The HET dark energy experiment"

Hirata C., Seljak U., 2003, MNRAS, 343, 459

Hu W., Haiman Z., 2003, PRD, 68, 3004

Huff E., Schulz A.E., White M., Schlegel D.J., Warren M.S., 2007, Astropart. Phys., 26, 351

Hui L., Gaztanaga E., LoVerde M., 2007, preprint [arxiv/0706.1071]

Hütsi G., A&A, 449, 891

Ibata, R.A., et al. 2003, MNRAS, 340, L21

Ida, S., & Lin, D. N. C. 2004a, ApJ, 604, 388

Ida, S., & Lin, D. N. C. 2004b, ApJ, 616, 567

Ida, S., & Lin, D. N. C. 2005, ApJ, 626, 1045

Irwin, M.J., et al. 2007, ApJ, 656, L13

Jura, M., Farihi, J., & Zuckerman, B. 2007, ApJ, 663, 1285

Juric, M., & Tremaine, S. 2007, ApJ, submitted (astro-ph/0703160)

Jurić, M., et al. 2007, ApJ (in press: astro-ph/0510520)

Kaiser, N. 1987, MNRAS, 227, 1

Kawka, A., & Vennes, S. 2006, ApJ, 643, 402

Kilic, M., et al. 2005, ApJ, 632, L115

Kilic, M., et al. 2006, AJ, 131, 582

Kogut A., et al., 2007, ApJ, in press.

Kollmeier, J.A., & Gould, A. 2007, ApJ, 664, 343

Koopmans, L. V. E., Treu, T., Bolton, A. S., Burles, S., & Moustakas, L. A. 2006, ApJ, 649, 599

Kornet, K., Bodenheimer, P., Różyczka, M., & Stepinski, T. F. 2005, A&A, 430, 1133

Köster, D., & Knist, S. 2006, A&A, 454, 951

Kravtsov, A.V., et al. 2004, ApJ, 609, 35

Lambert, D.L., Brown, J.A., Hinkle, K.H. & Johnson, H.R. 1984, ApJ, 284, 223

Law, D.R., Johnston, K.V., & Majewski, S.R. 2005, ApJ, 619, 807

Lépine, S., & Shara, M.M. 2005, AJ, 129, 1483

Liebert, J., et al. 2003, AJ, 126, 2521

Lissauer, J.J. & Stevenson, D.J. 2007, "Formation of Giant Planets", in Protostars and Planets V, B. Reipurth, D. Jewitt, & K. Keil, eds. (Tucson: University of Arizona Press)

Lomb, N.R. 1976, Ap&SS, 39, 447

Lovis, C., et al. 2006, Nature, 441, 305

Lupton, R. H., Gunn, J. E., Ivezić, Z., Knapp, G. R., Kent, S., & Yasuda, N. 2001, in ASP Conf. Ser. 238: Astronomical Data Analysis Software and Systems X, Vol. 10, 269

Majewski, S.R., et al. 2003, ApJ 599, 1082

Martínez-Delgado, D., Peñarrubia, J., Jurić, M., Alfaro, E.J., & Ivezić, Ž. 2007, ApJ, 660, 1264

McArthur, B.E., et al. 2004, ApJL, 614, 81

McDonald P., et al., 2000, ApJ, 543, 1

McDonald P., et al., 2005, ApJ, 635, 761

McDonald P., Eisenstein D., 2006, PRD, in press [astro-ph/0607122]

Meiksin A., White M., Peacock J.A., 1999, MNRAS, 304, 851

Meléndez, J., Barbuy, B., Bica, E., Zoccali, M., Ortolani, S., Renzini, A. & Hill, V. 2003, A&A, 411, 417

Merrill, K.M. & Ridgway, S.T. 1979, ARA&A, 17, 9

Mo, H. J., Mao, S., & White, S. D. M. 1998, MNRAS, 295, 319

Moore, B., Diemand, J., Madau, P., Zemp, M., & Stadel, J. 2006,MNRAS, 368, 563

Morrison, H.C., et al. 2001, AJ, 121, 283

Morrison, H.C., et al. 2003, AJ, 125, 2502

Munn, J.A., et al. 2004, AJ 127, 3034

Navarro, J.F., et al. 2004, MNRAS, 349, 1039

Navarro, J. F., & White, S. D. M. 1994, MNRAS, 267, 401

Navarro, J. F., & Steinmetz, M. 2000, ApJ, 528, 607

Newberg, H.J., et al. 2002, ApJ, 569, 245

Norris, J.E. et al. 2007, submitted to ApJ

Nishimichi T., et al., 2007, preprint [arxiv:0705.1589]

Olling, R. P., & Merrifield, M. R. 1998, MNRAS, 297, 943

Origlia, L., Valenti, E. & Rich, R.M. 2005, MNRAS, 356, 1276

Padmanabhan N., et al., 2007, MNRAS, 378, 852

Padmanabhan, N. et al. 2007, preprint (astro-ph/0703454)

Peebles P.J.E., Yu J.T., 1970, ApJ, 162, 815

Percival W., et al., 2007a, ApJ, 657, 51

Percival W., et al., 2007b, preprint [arxiv/0705.3323]

Perlmutter S., et al., 1999, ApJ, 517, 565

Perryman, M.A.C., et al. 2001, A&A, 369, 339

Preston, G.W., Shectman, S.A., & Beers, T.C. 1991, ApJ, 375, 121

Prugniel, Ph., Soubiran, C., Koleva, M. & Le Borgne, D. 2007, astro-ph/0703658

Ratra B., Peebles P.J.E., 1988, Phys. Rev. D37, 3406

Rich, R. M. & Origlia, L. 2005, ApJ, 634, 1293

Read, J.L., Pontzen, A.P., & Viel, M. 2006, MNRAS, 371, 885

Re Fiorentin, P., et al. 2007, A&A, 467, 1373

Riess A.G., et al., 1998, AJ, 116, 1009

Rivera, E.J., et al. 2005, ApJ, 634, 625

Robertson, B., Yoshida, N., Springel, V., & Hernquist, L. 2004, ApJ, 606, 32

Robertson, B., Bullock, J.S., Font, A.S., Johnston, K.V., & Hernquist, L. 2005, ApJ 632, 872

Rocha-Pinto, H. J., Majewski, S. R., Skrutskie, M. F., & Crane, J. D. 2003, ApJ, 594, L115

Sackett, P. D. 1997, ApJ, 483, 103

Sales, L.V., Navarro, J.F., Abadi, M.G., & Steinmetz, M, 2007, MNRAS, 379, 1464

Santos, N.C., Israelian, G., & Mayor, M. 2004, A&A 415, 1153

Scargle, J.D. 1982, ApJ, 263, 835

Schiavon, R.P., Barbuy, B. & Singh, P.D. 1997, ApJ, 484, 499

Seljak, U., et al. 2005, Phys. Rev. D71, 103515

Seo H.-J., Eisenstein D.J., 2003, ApJ, 598, 720

Seo H.-J., Eisenstein D.J., 2005, ApJ, 633, 575

Seo H.-J., Eisenstein D.J., 2007, preprint [astro-ph/0701079]

Sheldon E., et al., 2004, AJ, 127, 2544

Skrutskie, M.F. et al. 2006, AJ, 131, 1163

Smith R.E., Scoccimarro R., Sheth R.K., 2007, preprint [astro-ph/0703620]

Smith, V.V., Suntzeff, N.B., Cunha, K., Gallino, R., Busso, M., Lambert, D.K. & Straniero, O., 2000, AJ. 119, 1239

Springel V., et al., 2005, Nature, 435, 629

Stinson, G., Seth, A., Katz, N.S., Wadsley, J., Governato, F., & Quinn, T. 2006, MNRAS, 373, 1074

Sunyaev R., Zel'dovich Ya.B., 1970, Astrophysics & Space Science, 7, 3.

Tegmark, M. et al. 2004, ApJ, 606, 702

Tegmark M., et al., 2006, Phys. Rev. D74, 123507.

Thompson, R.I., Schnopper, H.W., Mitchell, R.I. & Johnson, H.L. 1969, ApJ, 158, L117

Tumlinson, J. 2006, ApJ, 641, 1

Udry, S., et al. 2006, A&A, 447, 361

Udry, S., Fischer, D., & Queloz, D. 2007, "A Decade of Radial-Velocity Discoveries in the Exoplanet Domain", in Protostars and Planets V, B. Reipurth, D. Jewitt, & K. Keil, eds. (Tucson: University of Arizona Press)

Valenti, J.A., & Fischer, D.A. 2005, ApJS, 159, 141

Vallinotto A., Dodelson S., Schimd C., Uzan J.-P., 2007, preprint [astro-ph/0702606]

Vivas, A. K., & Zinn, R. 2006, AJ, 132, 714

von Hippel, T., et al. 2007, ApJ, 662, 544

Wan, X.K., et al. 2006, Proc.SPIE, 6269, 88

Wetterich, C., 1988, Nucl. Phys., B302, 668

White M., 2003, in the Davis Meeting on Cosmic Inflation [astro-ph/0305474]

White M., 2005, Astropart. Phys., 24, 334

Wilhelm, R., et al. 1999, AJ, 117, 2329

Willman, B., et al. 2005a, AJ, 129, 2692

Willman, B., et al. 2005b, 2005, ApJ, 626, L85

Wu, Y., Murray, N.W., & Ramashi, J.M. 2007, ApJL, submitted (arXiv:0706.0732)

Wyse, R.F.G., et al. 2006, ApJ, 639, L13

Zehavi, I., Eisenstein, D. J., Nichol, R. C., Blanton, M. R., Hogg, D. W., Brinkmann, J., Loveday, J., Meiksin, A., Schneider, D. P., & Tegmark, M. 2005, ApJ, 621, 22

Zhao, B., & Ge, J. 2006, Proc.SPIE, 6269, 89

Zhao, G., et al. 2006, CJAA, 6, 265

Zhao, H., Spergel, D. N., & Rich, R. M. 1994, AJ, 108, 2154

Zhou, J.-L., Aarseth, S. J., Lin, D. N. C., & Nagasawa, M. 2005, ApJ, 631, L85

Zhou, J.-L., Lin, D. N. C., & Sun, Y.-S. 2007, ApJ, accepted (arXiv:0705.2164)

Zinn, R. 1993a, in "The Globular Cluster-Galaxy Connection", ASP Conf. Ser. 48, 38

Zinn, R. 1993b, in "The Globular Cluster-Galaxy Connection", ASP Conf. Ser. 48, 302

Zucker, D.B., et al. 2006a, ApJ, 643, L103

Zucker, D.B., et al. 2006b, ApJ, 650, L41