2014 Shaw Lecture Essay in Astronomy

Cosmic Surveys and the Composition of the Universe

Shaun Cole[1], Daniel J. Eisenstein[2], John A. Peacock[3]

1: Institute for Computational Cosmology, Durham University, UK

2: Harvard-Smithsonian Center for Astrophysics, Cambridge, USA

3: Institute for Astronomy, University of Edinburgh, UK

1. Introduction

The large-scale structure in the distribution of galaxies in the universe represents one of the great discoveries of modern cosmology. Three-dimensional maps from redshift surveys trace out a rich structure of galaxy clusters connected into superclusters by filaments and walls and surrounding interconnected voids, revealing structures of over 100 million light years in extent. This cosmic web of galaxies is not only intrinsically beautiful but also contains information that can be used to answer some of the major questions in cosmology:

1. What is the composition of the universe?
    - What is dark matter?
    - What is dark energy?
2. How did structure arise?

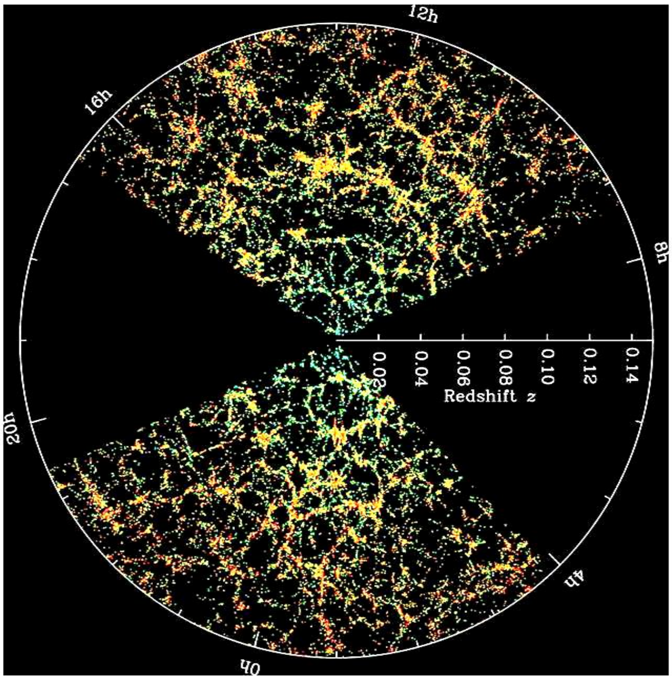

Figure 1: The large-scale cosmic web of galaxies as revealed by the Sloan Digital Sky Survey. We are at the centre of the diagram. Image credit: Michael Blanton and the SDSS Collaboration.

The key process that allows astronomers to build large 3-dimensional galaxy surveys of the local universe is the cosmic expansion that was discovered about 100 years ago. Between 1912 and 1923, Slipher showed that the light from galaxies is redshifted due to the Doppler effect. In 1929, Hubble found evidence that the redshift increased linearly with the distance to the galaxy. The relationship known as Hubble's law can be written as $v = H_0 d$, where d is the distance, v the velocity and H_0 is the current value of the Hubble expansion parameter. This is measured to be close to 70 km s^{-1} Mpc^{-1}, where 1 Mpc is roughly the typical distance between galaxies (3.26 million light years). This rate of expansion of the universe can change with time, which is a critical fact for the research we discuss here.

The utility of the Hubble expansion is that it can be exploited to construct large 3-dimensional maps of the surrounding galaxy distribution. Directly measuring galaxy distances is very hard as telescopes are unable to resolve individual stars in all but the nearest galaxies. The distance scale of Hubble's original diagram was incorrect due to bright clumps of star formation (HII regions) being mistaken for individual stars. But measuring the recession velocity of a galaxy is relatively straightforward. A telescope equipped with a spectrograph can measure the spectrum of light emitted by a galaxy and determine its redshift from the wavelength shift of standard emission lines and other spectral features. Between them, the SDSS-I and 2dFGRS surveys measured the redshifts of more than a million galaxies and so turned 2-dimensional catalogues of the angular positions of galaxies on the sky into a 3-dimensional distribution of galaxies using the redshift as the radial coordinate.

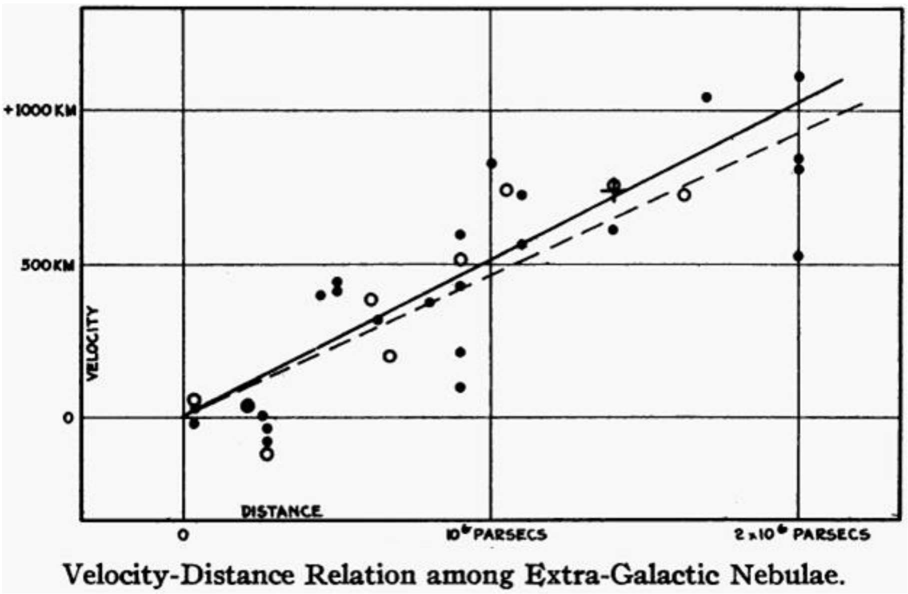

Figure 2: Hubble's original 1929 diagram revealing the linear relationship between the distances to galaxies and their recession velocities (Hubble 1929).

The Hubble expansion implies that the universe is expanding, galaxies are moving further apart from each other and the material in the universe is cooling. Conversely the universe must have been hotter and denser in the past: the so-called hot Big Bang. This model was provided with strong confirmation in 1964 when Arno Penzias and Robert Wilson detected the 3 Kelvin cosmic microwave background (CMB) − the thermal afterglow of the Big Bang. In the 1940s

George Gamow, Ralph Alpher and Robert Herman had studied Big Bang Nucleosynthesis (BBN) the nuclear fusion reactions that would happen during the very hot first few minutes of the universe. They deduced that light elements such as Helium would be formed during this time and for these elements to have abundances that matched observations the current temperature of the thermal radiation of the Big Bang would have to be a few degrees Kelvin. Hence the importance of the detection of the CMB and the 1992 confirmation by NASA's COBE satellite that it has an essentially pure black body spectrum.

Another important ingredient in the modern cosmological model is dark matter, which interacts gravitationally with ordinary matter but has no electromagnetic interactions. Evidence for the existence of dark matter first came in 1933 from the work of Fritz Zwicky, who found that galaxies in the Coma galaxy cluster were moving too fast to be bound into a stable system by the mutual gravitational attraction of their visible stars.

The expansion of the universe is governed by the balance between gravity, described by Einstein's theory of General Relativity, and the initial impetus of the Big Bang. If the only content of the universe were matter, then the mutual gravitational attraction of mass in the universe would act to decelerate the cosmic expansion. If the density were above a critical value dependent on the Hubble parameter, $\rho_{crit} = 3H_0^2/8\pi G$, where G is the gravitational constant, then the expansion would eventually be reversed and the universe would re-collapse, while if the density were lower the universe would expand forever but at an ever decreasing rate. Because of the importance of this critical density, we define the density parameter, Ω, which is density in units of the critical value: $\Omega = 8\pi G\rho/3H_0^2$. The density is also important in determining the geometry of the universe: if $\Omega > 1$, the universe would be closed in the same sense as the surface of a sphere: spacetime would be so strongly curved that a traveller proceeding in a straight line could return to their starting point. A universe with exactly $\Omega = 1$ would be termed 'flat'.

As it turns out, the universe contains more than just matter, and so one of the most important tasks for cosmology is to measure the Ω values corresponding to different constituents. We know that cosmology is complicated in this way because the expansion of the universe is accelerating, rather than decelerating as expected. This can be seen directly from Hubble's Law: $v = H\,d$. If we look to sufficiently large distances, we see the universe as it was in the past owing to the finite speed of light, so the change of H with time can be measured if we have data on the distance-redshift relation to moderately high redshift. The work described here represents one way in which such studies can be carried out. An alternative route is to use Type-Ia supernovae as standard candles (Shaw Prize in Astronomy 2006; Nobel Prize 2011), and the conclusions of these two routes confirmed each other in the late 1990s: the cosmic expansion is speeding up today. In the context of Einstein's General Relativity this can only be understood if the energy density of the universe is not dominated by normal (or dark) matter, but instead by some form of Dark Energy associated with the vacuum of space − i.e. a substance that does not clump but is spread uniformly.

Mathematically the simplest form of Dark Energy is equivalent to the cosmological constant, which Einstein introduced in 1917 in order to permit a stationary universe − since Einstein was unaware of Slipher's measurements showing that the galaxies were in motion. Thus Einstein introduced the cosmological constant to provide a repulsive force that balanced the large scale attractive gravitational force produced by the mass in the universe. Einstein withdrew his cosmological constant when he heard of the evidence for an expanding universe, but something akin to the cosmological constant is what is needed to produce the now observed accelerating expansion. The cosmological constant while being mathematically simple corresponds to a fixed unchanging vacuum energy density, which lacks physical motivation. Consequently many different physical dark energy models have been postulated in which the dark energy is related to a new physical field, motivated by the Higgs field in particle physics, which could have interactions with other components of the universe and evolve in a variety of different ways. The nature of dark energy is far from understood and we need better constraints on its properties. The way in which the dark energy and the other components of the overall energy density of the universe evolve over cosmic time affects how the expansion rate of the universe evolves. Hence by making accurate measurements of the past expansion history we can determine the composition of the universe and constrain the properties of the dark energy.

The Type Ia supernovae used to detect dark energy are good but not perfect standard candles. We do not fully understand why they have the luminosities they do, nor do we understand the extent to which their luminosities may be influenced by the evolving environments in which they form. Hence they are not perfect tools for the measurement of cosmic distances and it would be better if we had a cosmic ruler of known length. The phenomenon of Baryonic Acoustic Oscillations (BAO) provides such a cosmic ruler. Its origin in the physics of the early universe also provides a test of the model of the origin of structure and the properties of the matter components of the universe. Hence the detection of the BAO feature in the large-scale distribution of galaxies is an important step in determining the composition of the universe, and further measurements with it can provide more precise measurements of the properties of dark energy.

Origin and Growth of Fluctuations

The universe is not bland and uniform but instead contains complex structures from terrestrial scales, through the scales of the solar system and galaxy, out to the large structure of the cosmic web. The leading theory for the origin of the seeds of this structure is that it started with microscopic quantum fluctuations, which were stretched to macroscopic scales by an early epoch of very rapid expansion. The Inflationary Theory (for which Alan Guth, Andrei Linde and Alexei Starobinsky were awarded the Kavli Prize in 2014) postulates that in the first 10^{-34} seconds of the Big Bang the universe expanded exponentially by more than a factor of 10^{26} and in so doing stretched standard quantum fluctuations to scales that are now larger than our entire observable universe. These density fluctuations were initially of very low amplitude but over cosmic time they are amplified by gravitational instability. The gravitational attraction of mass in slightly overdense regions is slightly greater and so gravity causes these regions to expand less rapidly than their surroundings and so become more overdense. Conversely underdense regions expand more rapidly and become more underdense.

For the first 400,000 years after the Big Bang, the normal or baryonic matter in the universe is in the form of a highly ionized photon-rich plasma. During this time, the pressure in this plasma plays a role in the evolution of density fluctuations in the baryons. However the dark matter has no collisions with the photons or baryons and so does not feel this pressure. After 400,000 years the temperature drops below 3000 Kelvin and the point is reached where the ionized plasma can recombine to form a gas of neutral hydrogen atoms. By this time the density fluctuations have grown to a typical amplitude of 1 part in 100,000. The CMB provides us with a glimpse of the universe at this epoch, as after recombination the universe is transparent and CMB photons travel to us from this last scattering surface undisturbed. Maps of these temperature fluctuations across the microwave sky made by the COBE, WMAP (2010 Shaw Prize) and Planck satellites reveal a spectrum of Gaussian temperature fluctuations in good accord with Inflation.

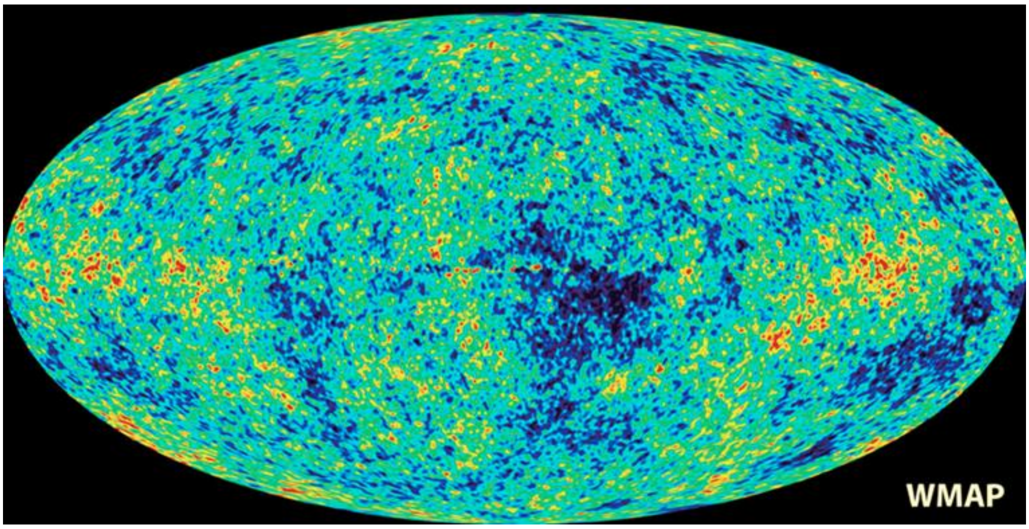

Figure 3: The WMAP all sky map of fluctuations in the temperature of the cosmic microwave background. The differences in temperature between the hot and cold regions are just a few parts in a hundred thousand. Foreground emission from the Milky Way has been subtracted, to yield a picture of the fluctuations in the distribution of matter as they were about 400,000 years after the big bang. In inflationary cosmology, these structures were created by amplified quantum fluctuations. In any case, they are the visible seeds of the large-scale structure in the present-day universe. Image credit: Bennett et al. 2013.

In the intervening 13 billion years, gravitational instability has continued to amplify these density fluctuations to the point where they have become non-linear — density fluctuations of order unity and greater. This non-linear evolution can be followed using large computer simulations which represent the mass distribution in the universe by a collection of N discrete particles or bodies. N-body simulations have been run in which more than 100 billion particles are used to represent the mass distribution in a large representative volume of the universe. These simulations reveal that the action of gravitational instability operating on the linear fluctuations predicted by Inflation and seen in the CMB is a cosmic web of filaments and walls in which structure forms hierarchically. First dense clumps of matter form, then these clumps stream along walls and filaments and merge to form larger structures.

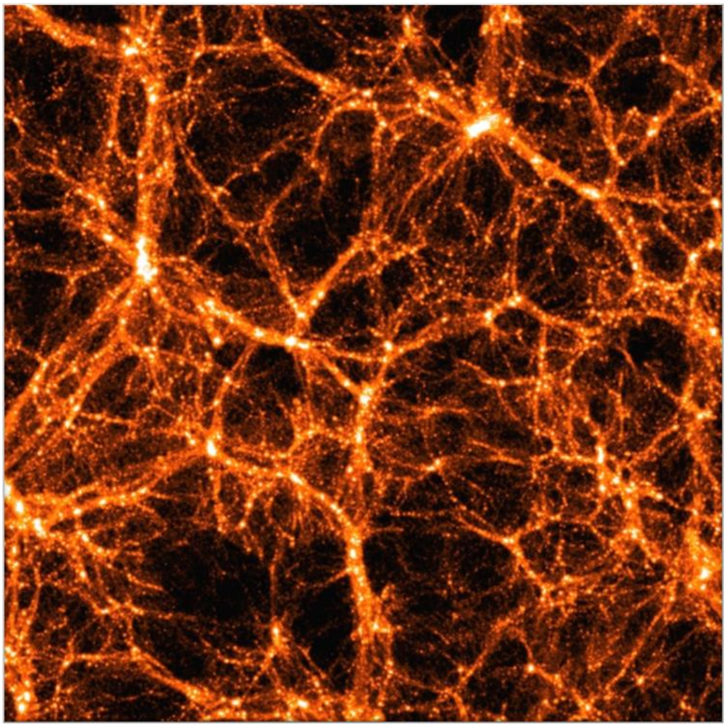

Figure 4: A slice through an N-body simulation of a 300 million light year (100 Mpc) cube of the universe. The colour scale shows the distribution of dark matter which is clumped into dense concentrations connected by filamentary structures, surrounding nearly empty voids. This structure is the outcome of gravitational attraction enhancing and modifying the small initial fluctuations seen in the CMB. But although the small-scale structure is complex, the residual fluctuations on large scales can be simply related to those in the initial conditions.

It is this whole model of the formation of structure from quantum fluctuations processed by gravity and pressure in the early universe and further amplified by gravity in the later non-linear regime that we wish to test by mapping and quantifying the large-scale galaxy distribution.

2. Observing Large-Scale Structure

The traditional method of measuring a galaxy redshift (that used by Slipher) is to use a large telescope to gather the light from a single galaxy and to focus this onto the input slit of a spectrograph. Rather like a prism, this splits the light into to its component wavelengths (colours). As galaxy spectra contain characteristic emission lines (formed in ionized gas around hot young stars) and absorption lines (formed in the atmospheres of cool giant stars), the measured shift of these lines from their standard rest frame values can be used to measure the galaxy redshift. As the measurement of each galaxy requires observing for a significant time to build up a sufficiently high signal-to-noise spectrum, this is a long laborious process; Consequently, even as recently as 1980 only of order 1000 galaxies had measured redshifts.

The first systematic galaxy redshift survey, CfA-I, was undertaken between 1977 and 1982 and consisted of the 2400 brightest galaxies at high galactic latitude (Davis et al 1982). The next leap in galaxy redshift surveys came with the Las Campanas Redshift Survey (LCRS; Shectman et al 1996) which between 1991 and 1996 measured 26,000 redshifts. This order of magnitude increase was achieved by using a multiplexed spectrograph fed by optical fibres capable of simultaneously measuring 50 galaxy spectra. However this multiplex capability was soon boosted by another order of magnitude by the 2dF and SDSS spectrographs.

The 2-degree Field Galaxy Redshift Survey

The 2dF facility was selected for funding by the Anglo-Australian Observatory in 1990 and became fully operational on the 3.9-m Anglo-Australian telescope in 1995. Its novel design consisted of two sets of 400 optical fibres spread over a 2-degree field of view and each connected to small prisms which collect the light from an individual galaxy and feed it into its optical fibre. The other ends of the optical fibres are lined up to feed their light into two optical spectrographs, which simultaneously record the 400 galaxy spectra during a 45-minute exposure. While observations are taking place with one set of 400 fibres a pre-programmed robot positions the other set of 400 fibres on a second focal plate. When the first set of observations is complete, the central barrel shaped portion of the 2dF instrument tumbles to place the second focal plate in the focal plane ready to collect the light from the next set of 400 galaxies. In this way, the 2dF facility is capable of measuring 400 galaxy redshifts every 45 minutes.

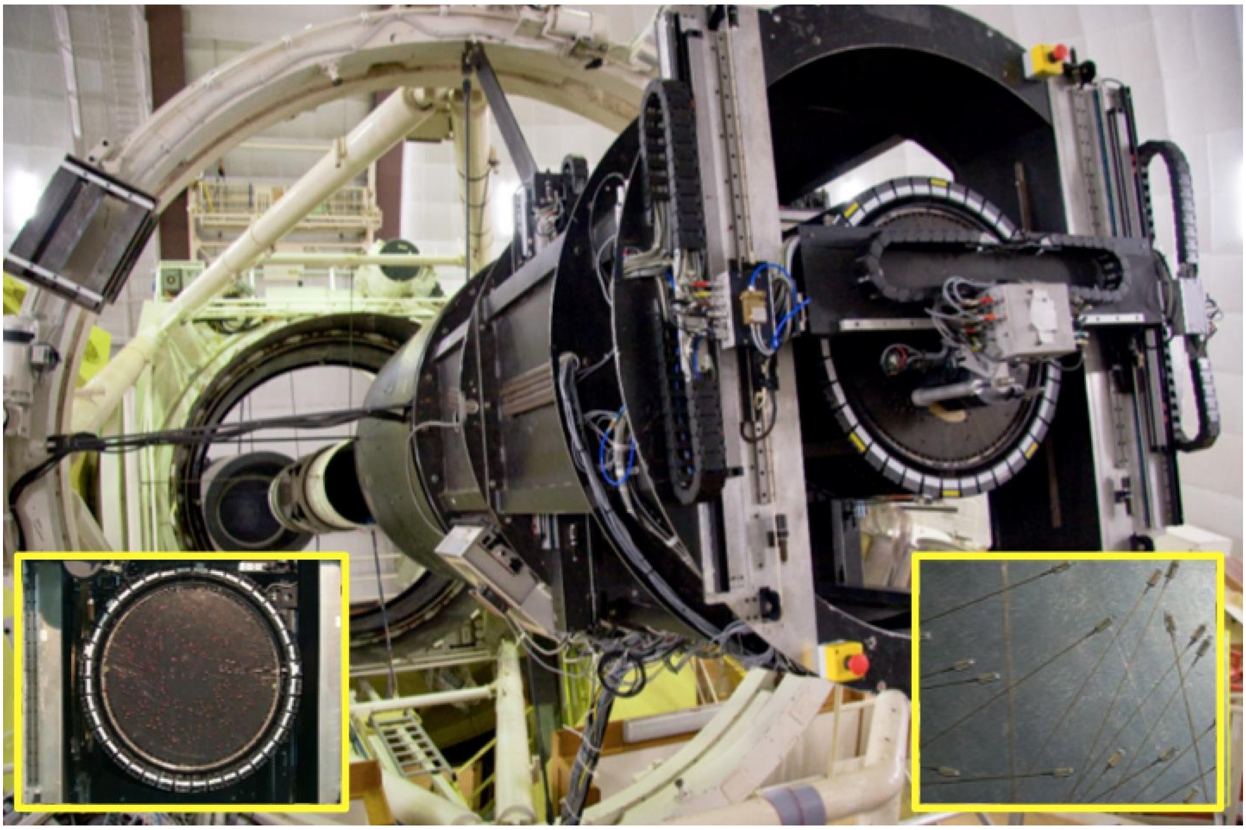

Figure 5: The 2dF instrument mounted on the 3.9-m Anglo-Australian telescope. The primary mirror is visible on the left and the robot positioning fibres on the right. The two insets show the distribution of fibres on the whole focal plane and in close up. Image courtesy of the Anglo-Australian Observatory.

The capabilities of the 2dF facility were one key ingredient that enabled the 2dF team to carry out the 2dF galaxy redshift survey. The other key ingredient was the input galaxy catalogue based on the APM galaxy survey led by Steve Maddox and George Efstathiou (Maddox et al 1990). This is a homogeneous complete two-dimensional catalogue of galaxies constructed from digital scans of wide field photographic plates taken with the UK Schmidt telescope. Together these allowed the 30 strong Anglo-Australian 2dF team to measure over 220,000 galaxy redshifts during 300 nights of prime observing time between 1995 and 2003 (Colless et al 2001). These enabled portions of the flat 2-dimensional APM galaxy catalogue to be made into 3-dimensional galaxy catalogues, revealing the cosmic web and allowing precise quantification of the 3-dimensional galaxy distribution.

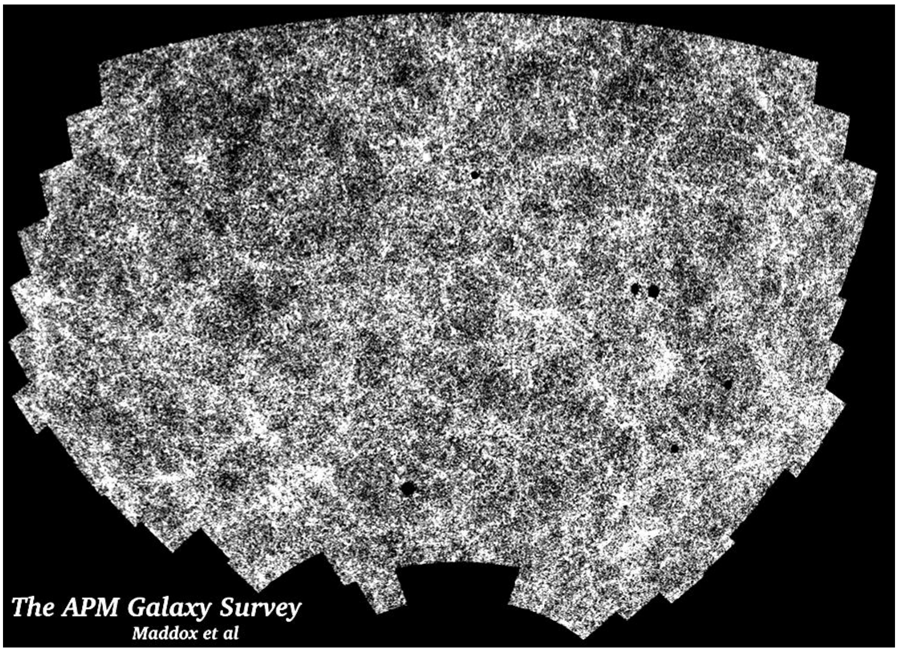

Figure 6: The distribution of galaxies on the sky in the APM galaxy survey. The image depicts over 2 million galaxies in a region 100 degrees across centred towards the south pole of our Milky Way galaxy. Bright regions indicate more galaxies. Dark ellipses indicate where bright stars have been masked out. There are weak indications of filamentary structure in the galaxy distribution, but these can only be seen clearly once we have the 3rd dimension that comes from spectroscopy. Image credit: Steve Maddox

The scientific legacy of the 2dF galaxy redshift survey has been broad and far-reaching. There have been over 100 scientific papers published based on 2dFGRS data covering everything from the properties of radio galaxies to setting a constraint on the mass of neutrinos. Forty of these papers are by the 2dF team, and the others were made possible by the prompt release of all the 2dFGRS data in a public database (Colless et al. 2001).

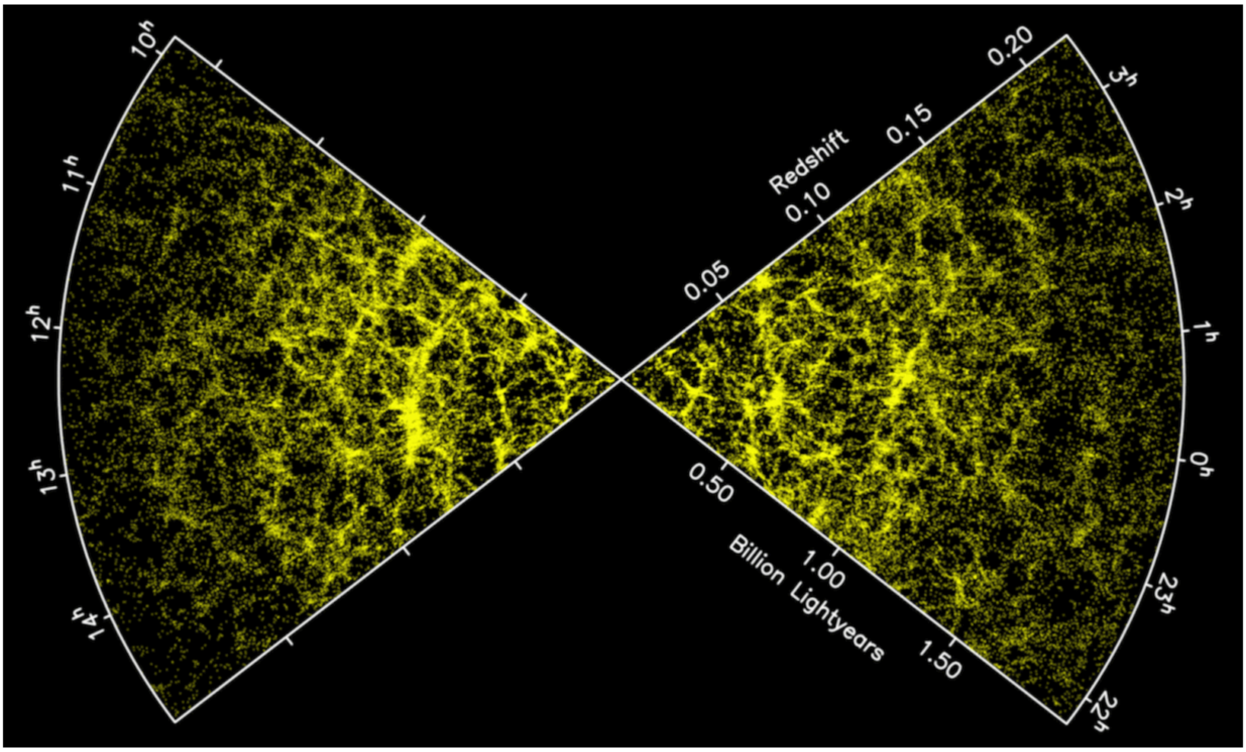

Figure 7: Cone plots showing slices through the 3-dimensional distribution of galaxies in the 2dF Galaxy Redshift Survey. The observer is at the centre of the plot. Data have been stacked in the thin (~4 degree) direction to create a two-dimensional image. One can see the clumpy filamentary structure of the cosmic web, bearing a striking resemblance to the structures seen in theoretical simulations of gravitational collapse, such as Figure 4. It is also worth noting that the two slices correspond to strips in nearly opposite directions in the sky, and yet the statistical character of the fluctuations appear very similar on each side. This gives reassurance that our surveys are deep enough to yield a representative sample of the universe. Image credit: The 2dFGRS Team.

The Sloan Digital Sky Survey

The SDSS project began in 1990 with the goal of leveraging the revolution in digital detectors to make the largest yet map of the universe (York et al. 2000). Using a custom-built wide-field 2.5-m telescope at Apache Point Observatory, the SDSS mapped the sky using two powerful instruments. The first was a wide-field digital camera, the largest of its kind at the time, capable of producing deep and wide 5-colour images of the sky. With this camera, the SDSS imaged one third of the sky, detecting over 200 million galaxies and 250 million stars. The second was a multi-fibre spectrograph, which could take spectra of 640 objects simultaneously from a 3-degree diameter field. Both instruments had very high throughput that, combined with the wide field of view of the telescope, allowed astronomers to map the sky at an unprecedented rate. By 2009, the SDSS had taken spectra of over 1.5 million objects, including nearly a million galaxies and quasars.

But beyond data volume, the true revolution of the SDSS was the digital and systematically uniform nature of the data, which permitted exceptionally versatile and well-characterized processing of images and spectra. Because of this, the data set has supported highly refined statistical applications, from searches for 1-in-a-million rare objects to subtle cosmological correlations arising in samples of millions of galaxies. The full data, from raw images to science-grade catalogues, have been released to the worldwide astronomy community and the public in a series of flexible and feature-rich databases. To date, over 4000 refereed papers have been written from SDSS data, touching nearly all fields of astrophysics, from the nearest asteroids to the most distant quasars.

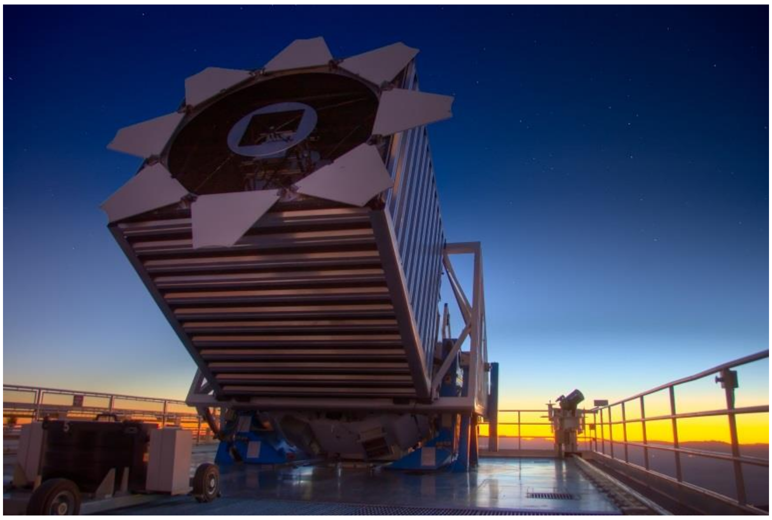

Figure 8: The Sloan Foundation Telescope, on which the SDSS has been performed. The 2.5-m telescope at Apache Point Observatory at 9200 feet elevation in New Mexico, US, is remarkable for its wide field of view, 3 degrees in diameter. Unlike many telescopes, it does not sit in a dome but instead is protected from the wind by an enshrouding baffle (the slats visible here) that moves with the telescope. Eight petals at the top of the baffle can be folded into the view of the telescope to provide a screen against which to project light from calibration lamps. Not seen here is the shed that protects the telescope from rain, snow, and sun; this rolls off at night to reveal the telescope for use. Photo credit: David Kirkby.

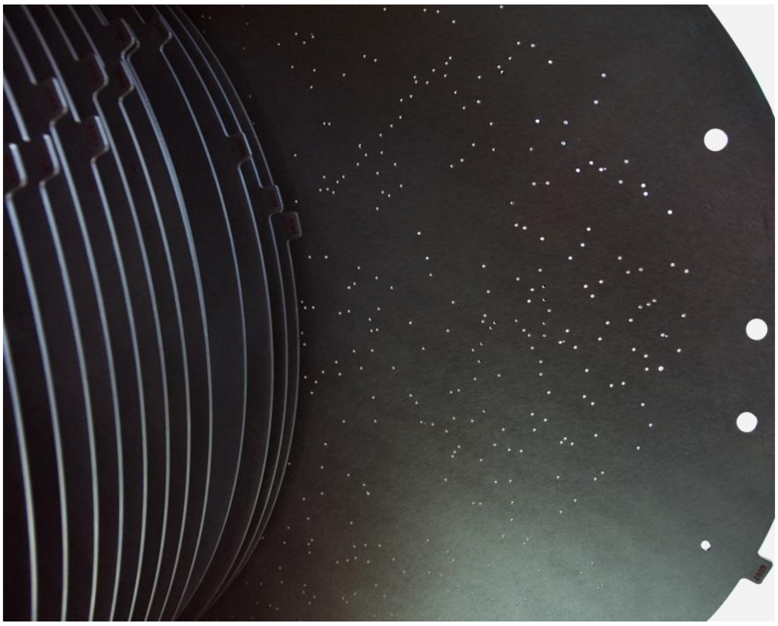

Figure 9: The SDSS spectrograph positions its fibres mechanically using holes drilled into thin aluminium plates. Each hole receives a hand-plugged optical fibre. The plate is placed at the focal plane of the telescope, so that the light from the desired objects falls

into the fibre. This requires the holes to be drilled with 10 μm precision over the 60 cm field of view. Over 5000 plates have been observed in the past 15 years. Photo credit: David Kirkby.

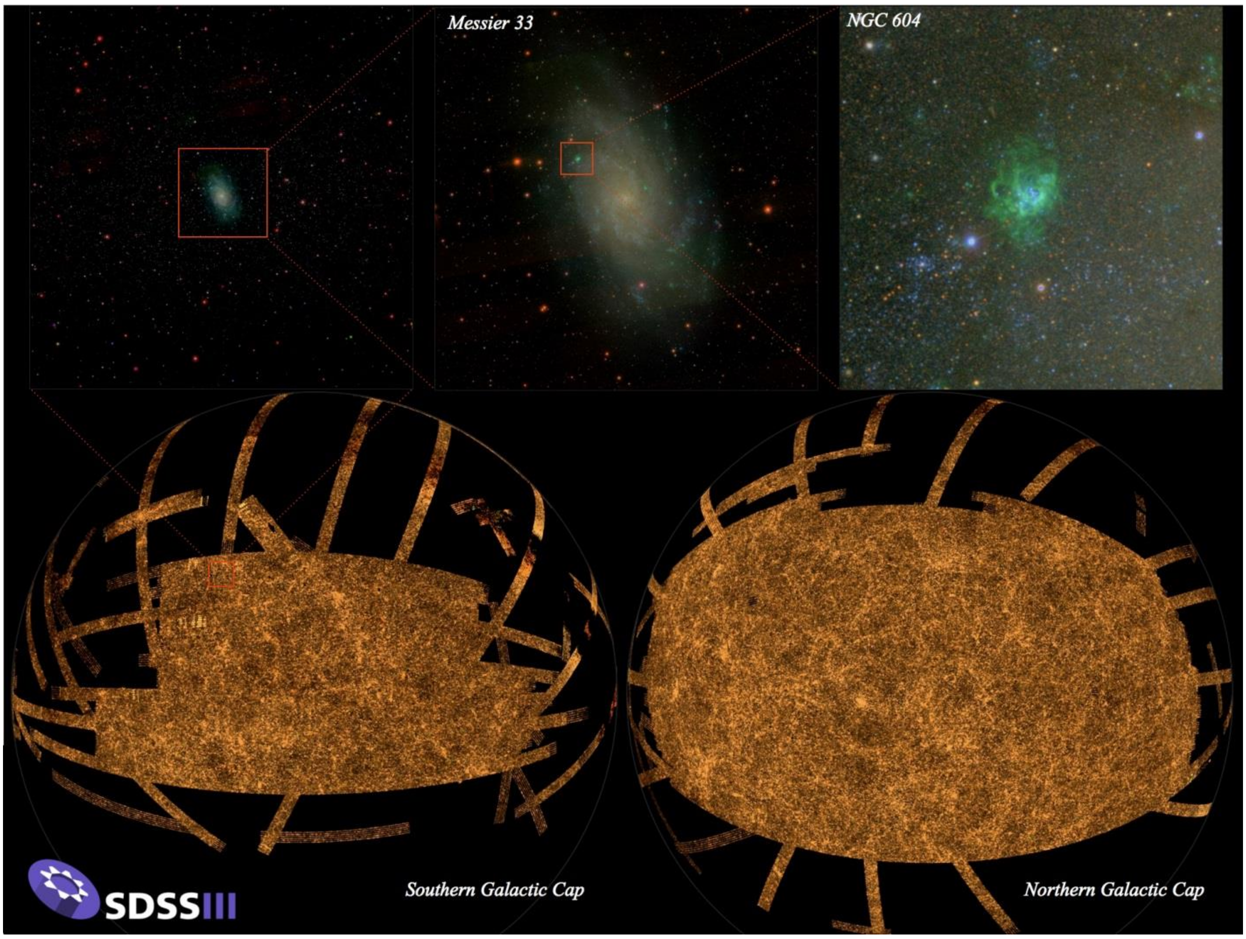

Figure 10: Composite image of the extra-galactic sky from the SDSS. The bottom row contains the hemispheric projections looking out of the Galactic disk toward the south and the north. Stars have been removed, and the light of the galaxies combined. The resulting pattern is due to the large-scale structure of the universe. The top row shows a zoom-in on a portion of the nearby galaxy Messier 33. This reveals the tremendous resolution of the image, over a trillion pixels, each digitally recorded in 5 bandpasses from the near-ultraviolet to the near-infrared. Image credit: Michael Blanton and the SDSS Collaboration.

The project received funding from the Alfred P. Sloan Foundation in 1992, with additional funding from the U.S. National Science Foundation, U.S. Department of Energy, several international funding agencies, and many member institutions. First imaging data arrived in 1998, with spectroscopic data coming a year later. The survey officially started in 2000. A second collaborative phase, SDSS-II, started in 2005 and completed the three-dimensional mapping of 8000 deg^2 (20% of the sky) for its extra-galactic survey. SDSS-II also conducted extensive studies of the Milky Way and conducted repeat imaging to measure dark energy with Type Ia supernovae.

A third collaborative phase, SDSS-III, began in 2008, with programs aimed at cosmology, the Milky Way, and extra-solar planets (Eisenstein et al. 2011). Most relevant to this essay, the Baryon Oscillation Spectroscopic Survey (BOSS) is the largest program in SDSS-III. With a substantial upgrade to the SDSS spectrograph that expands its grasp to 1000 simultaneous objects, BOSS has produced a new spectroscopic map of the universe, covering 25% of the sky and including targets a factor of 10 fainter than the original SDSS. The final maps from BOSS include nearly 2 million galaxies and quasars and provide a view of the large-scale structure of the universe 7 times more powerful than the previous SDSS. The SDSS-III collaboration includes over 800 scientists from more than 50 universities and research institutes worldwide. With SDSS-III, the Sloan Telescope facility has now acquired spectra of 2.5 million extragalactic objects, half of which have redshift above 0.5.

With the successful completion of SDSS-III, a fourth collaborative phase, SDSS-IV, began in summer 2014. SDSS-IV will continue to expand our survey of the sky, gathering yet higher redshift galaxies, dissecting nearby galaxies, and exploring new regions of the Milky Way. As we approach the 25^{th} anniversary of the initiation of the 2dF and SDSS projects, it is remarkable to see the continued success of one of the primary lessons of 2dF and SDSS: that a collaboration of diverse scientists can pursue many goals simultaneously through a coordinated systematic sky survey.

We want to close this section by expressing our gratitude to the 2dF and SDSS collaborations. These projects were much more than instrumentation suites and data sets. Many talented people came together to design, construct, operate, analyse, and document these surveys. The collaborations became experts not only in how to pursue the sky surveys, but also in the nuances of the science applications of wide-field data sets. We all have experienced these projects to be vibrant scientific communities that have widely shaped our own study of the universe, and it has been our great privilege both professionally and personally to be a part of them.

3. Quantifying Large-Scale Clustering and Motion

The maps of the cosmological large-scale structure are beautiful images, but they risk distracting us with irrelevant detail. As far as we know, there is no fundamental significance to the exact placement of particular cosmic structures, and the distinction between a supercluster and a void may be no more meaningful than the difference between the crest and trough of a wave on the ocean. In other words, the fluctuations in cosmic density are a kind of random 'noise', where the density at a given point in space might with equal probability be above or below the average. We are used to trying to discern a signal in the presence of noise, such as a conversation over a particularly crackly telephone connection. But what is less familiar from everyday life is that the noise itself can contain information. Revealing this information requires us to take a statistical approach and average the properties of the density fluctuations over a large volume.

Clustering statistics

A useful way to think about the properties of large-scale structure is via Fourier analysis, where we imagine expressing the variation of density with position as a sum of many wave 'modes': where the density varies in an oscillating way as we move in a given direction, but where it is constant on surfaces perpendicular to the direction in which the wave runs. We are familiar with this on the surface of the sea: sometimes there is a single obvious dominant wave moving in a given direction, but very often there is no clear pattern, meaning that many waves moving in different directions are present simultaneously.

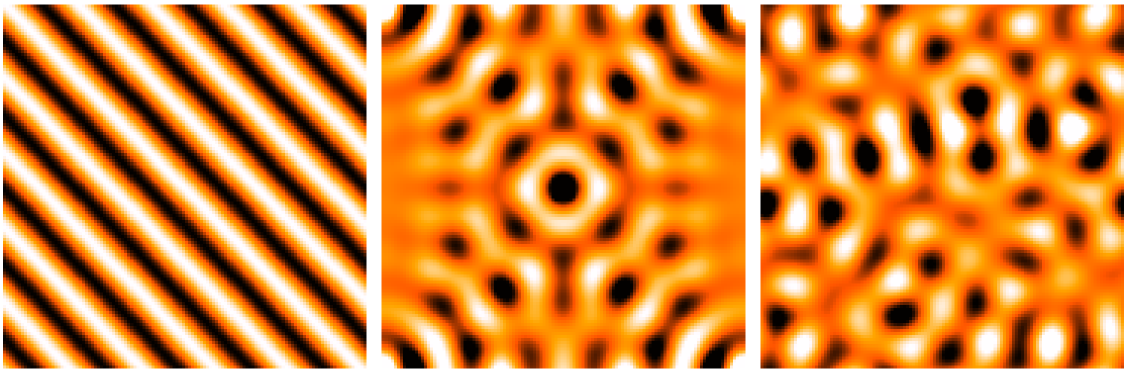

Figure 11: Illustrating how density fluctuations can be a Fourier superposition of waves. The first panel shows a single wave; the second adds together 13 similar waves running in different directions. The regularity of this pattern is destroyed when the waves are given a random shift in phase.

Figure 11 shows that adding the effects of many waves can give a regular pattern — but this regularity disappears if the waves are given a random phase shift, i.e. moved randomly along their direction of propagation. Unlike waves on water, we are not thinking of waves that physically travel: this is just a convenient mathematical way of representing something like the fluctuating density field. But given the density field, there is a unique way to extract the amplitude of each mode. These Fourier modes are characterised by a wavelength, λ, although it is more common to use the wave number, $k = 2\pi/\lambda$. By considering waves of different wavelengths, we can give a precise meaning to the intuitive idea that the density field contains structure on a range of scales.

The power spectrum is the tool that tells us the balance between the contributions from different scales: it simply means the square of the amplitude of the mode. Normally we are interested in the contribution to the density fluctuation from modes in some range of scales; specifically, the contribution to the fractional variance in density from a unit range of log(k). We will denote this by $\Delta^2(k)$.

A common alternative to the power spectrum is the correlation function, which encodes the characteristic sizes of cosmological structures such as superclusters and voids. Here we think about density fluctuations changing the probability of finding a neighbour to a given galaxy or matter particle. In the absence of clustering, this probability would be uniform at all radii, r. But owing to clustering, the probability is multiplied by a factor $1+\xi(r)$, where ξ is the correlation function. This tends to be large and positive at small separations, tending to zero at large separations to reflect the fact that the universe is uniform in the mean (although it can be slightly negative). Both the power spectrum and correlation function hold the same information. For example, if the universe was constructed out of clumps of size *L*, then the correlation function would drop to zero at roughly $r > L$, while the power spectrum would drop to zero at roughly $k < 1/L$.

Weighing the universe with the power spectrum

The power spectrum measured in the 2dFGRS is shown in Figure 12. This shows a number of features. The simplest is that the measurements increase as we go to smaller scales (larger *k*): the universe is lumpy on small scales, but becomes increasingly smooth as we view it with a coarse resolution (the contribution of low-*k* modes).

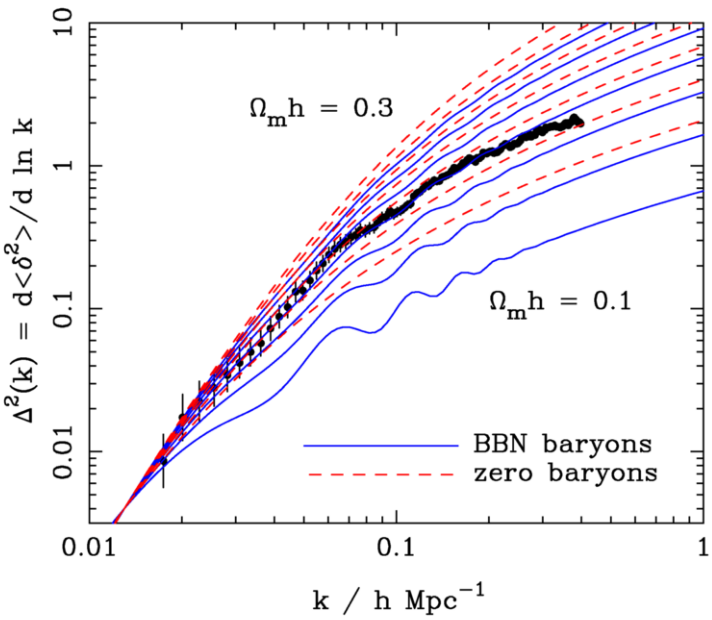

Figure 12: The 2dFGRS galaxy power spectrum: fractional variance in density per log scale, versus wavenumber. The data are contrasted with theoretical models in which the matter density of the universe is varied (data from Cole et al. 2005).

It is important that we can probe into the regime where $\Delta^2(k) \ll 1$, since this represents small fluctuations in density that are 'linear': the modes have evolved independently since they were created, so that they retain a clean picture of the initial conditions. It is also important to reach this regime of small fluctuations so that we can be confident that our surveys are large enough to constitute a representative sample of the whole universe.

The data are contrasted with theoretical models, which combine two ingredients. The initial inhomogeneities seeded (perhaps) by inflation would rise as a pure power law (straight line on this plot). But the fluctuations we measure today are reduced compared to these initial values on small scales, giving the spectrum a characteristic curvature. This happens because of the effects of pressure, which oppose those of gravity. For the component of the universe that is in the form of normal 'baryonic' matter, this pressure effect causes primordial sound waves, Baryon Acoustic Oscillations, which we discuss in more detail below. But such features are hard to see because dark matter is about 5 times more abundant than baryonic matter (this is one of the main ways we know that dark matter exists). For dark matter, the effect is simply that the fluctuations grow more slowly on small scales.

At a given time, the largest scale that can be affected in this way is the horizon distance: *ct*, where *c* is the speed of light. So the power spectrum gains a bend that moves to larger scales as the universe ages. But eventually, this process ceases, since the effects of pressure can only overcome gravity while the density of the universe is dominated by radiation (it is the radiation that supplies the pressure). So the large-scale structure is imprinted with the time of matter-radiation equality. This time is earlier for a high-density universe, so we can use this to weigh the universe. The break scale depends on the matter density parameter, $\Omega_m = \rho_m / \rho_{crit}$, and also the Hubble constant ($h = H_0$ / 100 km s^{-1} Mpc^{-1}). Since external data show that h is close to 0.7, this yields a matter density of Ω_m close to 0.3.

Does this mean that the universe is of below critical density? No, as we are only counting matter that can clump. In fact, we believe that the remaining 70% of the critical density is supplied by uniform dark energy, but this requires separate evidence.

Redshift-space distortions

All the discussion so far has treated the 3D information from the redshift as perfect. But this is not the case. The very existence of fluctuations in density means that there must also be 'peculiar velocities', in which galaxies move at a different speed than the one given by Hubble's law: for the density of a region to grow with time, galaxies must move towards each other. This means that the observed redshift is a combination of the Hubble expansion velocity and the Doppler shift from the radial component of the peculiar velocity. Our beautiful maps of the large-scale structure are therefore distorted, as illustrated in Figure 13.

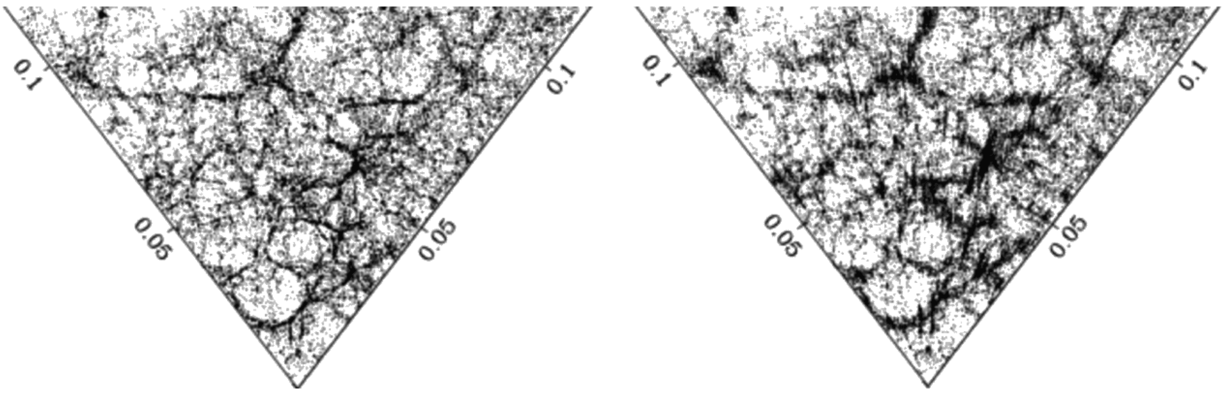

Figure 13: Simulations of the low-redshift portion of the 2dFGRS. On the left is what we would see in real space, with exact 3D distances. On the right, we see redshift space: whole superclusters are shifted in apparent position as a result of their mean velocity with respect to uniform expansion. But the random orbital velocities inside dense groups of galaxies causes smearing along the line of sight: the so-call fingers of God.

At first sight, this sounds like an ugly complication, but it turns out to be useful: the amplitudes of the peculiar velocities are interesting to know, since they give an alternative means of weighing the universe. To reveal these distortions statistically, consider the redshift-space correlation function measured from the 2dFGRS, as shown in Figure 14. Here, we are plotting the excess probability over random of one galaxy finding a neighbour. In the absence of peculiar velocities, this would just be a function of radius, r. But here, we measure separately the radial (π) and transverse (σ) components of this separation: and it can be seen that the contours of ξ are far from being circularly symmetric. Two distinct effects are apparent, first analysed thoroughly by Nick Kaiser in 1987. At small separations, there is a radial stretching, known as the 'finger of God' distortion: this represents random orbital speeds of galaxies moving together within the same group.

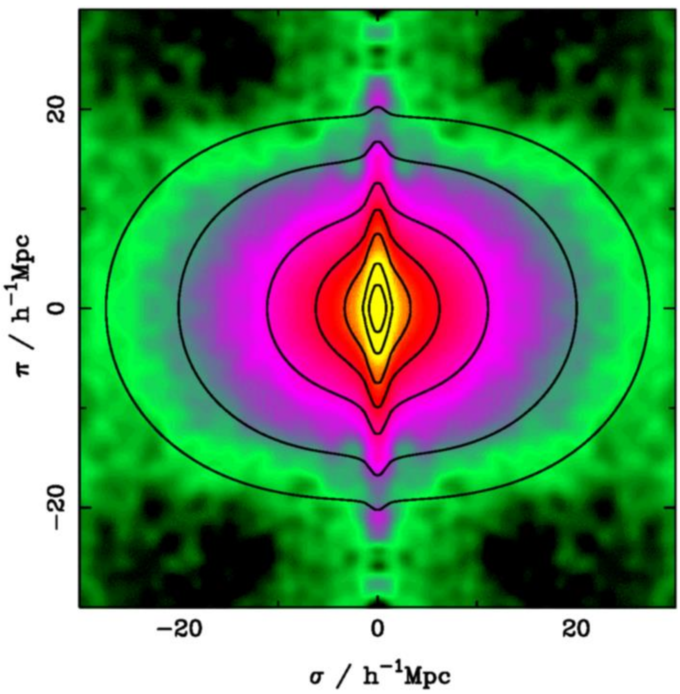

Figure 14: The redshift-space correlation function in 2dFGRS: the excess probability for a galaxy to find a neighbour, dissected according to transverse and radial separations. The latter is affected by peculiar velocities in two ways. At small separations, random orbital velocities within the same group cause the 'finger of God' radial stretching. But at large separations we see a flattening of the contours, reflecting the coherent infall of matter onto superclusters. The coloured data match extremely well with the contours of a theoretical model (figure from Peacock et al. 2001).

But of greater interest is the flattening of the correlation contours at large distances. This represents the velocities corresponding to modes that are still in the linear regime: matter falling together from large distances towards a supercluster. Objects on the far side of the supercluster have larger redshifts, but they fall towards us: this reduces the redshift we see and decreases the apparent distance, hence causing a squashing along the line of sight. The amount of flattening depends on the velocities, which are lower for a low-density universe.

Although detections of this 'Kaiser effect' had been made prior to the 2dFGRS, the full structure of the redshift-space correlations, as shown in Figure 14, had never been seen with the level of precision needed to appreciate the excellent match to theory. The observed degree of distortion depends on the matter density, and a value of Ω_m close to 0.3 was preferred, giving independent validation of the conclusion from the large-scale shape of the power spectrum. Once again, however, this measurement tells us nothing about the contribution of dark energy.

4. Baryon Acoustic Oscillations and Dark Energy

Prior to recombination, the photon-rich plasma has a high pressure and sound speed. When the photons dominate both the pressure and energy density of the plasma, the sound speed is 58% of the speed of light ($c_s = c/3^{1/2}$). Regions in the plasma that are overdense are also over-pressured. These pressure differences produce a force on the plasma that acts against gravity. The result is that perturbations in the baryons with wavelengths shorter than the distance sound can propagate since the Big Bang, oscillate as acoustic waves, rather like standing waves on a guitar string. In contrast, the dark matter, which does not interact with the photons, does not feel this pressure. Perturbations in the dark matter density simply grow under the influence of gravity. This results in the density waves in the baryons becoming progressively out of phase with those in the dark matter.

This continues until recombination, when the photons decouple from the baryons, resulting in a dramatic drop in the pressure and sound speed. The perturbations in the baryons cease oscillating. For wavelengths longer than the distance sound has travelled since the Big Bang, the sound horizon, the baryon fluctuations are still in phase with the dark matter as these perturbations have not begun to oscillate. However on smaller scales, the baryon perturbations are progressively more out of phase with those of the dark matter. Consequently the total spectrum of dark plus baryonic matter fluctuations exhibits wiggles on the scales where the two sets of fluctuations come in and out of phase (Peebles & Yu 1970, Sunyaev & Zel'dovich 1970).

The amplitude of these baryon wiggles depends on the fraction of matter that is baryonic. If the baryons were only a trace fraction, the wiggles would be very small and the total spectrum of density fluctuations would be dominated by the smooth spectrum of the dark matter. Therefore, detecting and measuring the amplitude of these wiggles provides a measurement of ratio of mass density of baryons and dark matter in the universe.

The characteristic scale of the Baryon Acoustic Oscillations (BAOs) — the length of our cosmic ruler — is simply set by the length of the sound horizon at recombination, i.e. how far sound propagates between the Big Bang and the epoch of recombination. This can be calculated accurately from first principles as the physics involved is all in the linear regime. The epoch of recombination, when the CMB radiation last interacted with matter, was a redshift of around z=1100 and an age of around 400,000 years. These figures are set almost entirely by the measured CMB radiation density, with only a very weak dependence on other quantities such as the mass density. Hence provided they can be detected, the BAOs are a superb cosmic standard ruler for carrying our geometrical measurements of the evolution of the size of the universe.

Since recombination, the universe has expanded and cooled by a factor of a thousand. The density fluctuations that have grown during this time to form the cosmic web were seeded by the combined sum of the baryonic and dark matter fluctuations. We therefore expect the length scale of the BAO cosmic ruler, or equivalently the baryon wiggles, to be encoded in the distribution of galaxies which trace out the cosmic web. The present-day length of the BAO scale is 500 million light years (150 Mpc) and the amplitude the BAO fluctuations are small, so this a very challenging feature to detect.

The 3-dimensional 2dFGRS and SDSS-I galaxy catalogues are the starting point for the analysis required to extract the BAO signal. They were the first surveys to cover a sufficiently large volume of space to make it feasible to detect BAOs.

The 2dFGRS analysis essentially involves modelling the properties of galaxies and the selection criteria of the survey in order to predict the smooth distribution of galaxies one would have seen if galaxies were just uniformly distributed in space. This then allows one to define the overdensity of galaxies as a function of position in the survey and use Fourier techniques to compute the spectrum of galaxy density fluctuations on very large scales. In reality, most of the work involved is in testing the assumptions made in modelling and calibrating the survey to ensure that systematic errors do not unduly influence the measurement.

In the SDSS-I analysis, a more direct way of detecting the length scale of the BAO feature was used. Here we calculated as a function of separation the auto-correlation function of the fluctuations in the galaxy density. In this statistic, the expected BAO feature is a single peak at a separation of 500 million light years. This morphology can be understood as follows: a region left overdense by inflation will be overdense in dark matter, but it will also generate a sound wave that travels outwards from the initial region. This wave travels 500 million light years before stalling at recombination, creating a small overdensity in the gas at that radius. Both the initial overdensity and the echo at 500 million light years favour conditions for forming galaxies. When we analyse the full map, galaxies are slightly more likely to be separated by 500 million light years than by 400 or 600 million light years.

In January 2005, SDSS and 2dF teams jointly announced the detection of the BAO feature in their galaxy clustering analysis. These measurements strongly confirmed the whole Big Bang model of structure formation and constrained the baryonic fraction of matter to be approximately 18%. The detection of the BAO provides a clear connection between the structure seen in the anisotropies of the microwave background 400,000 years after the Big Bang and that seen in the clustering of galaxies 10 billion years later (see Figure 15).

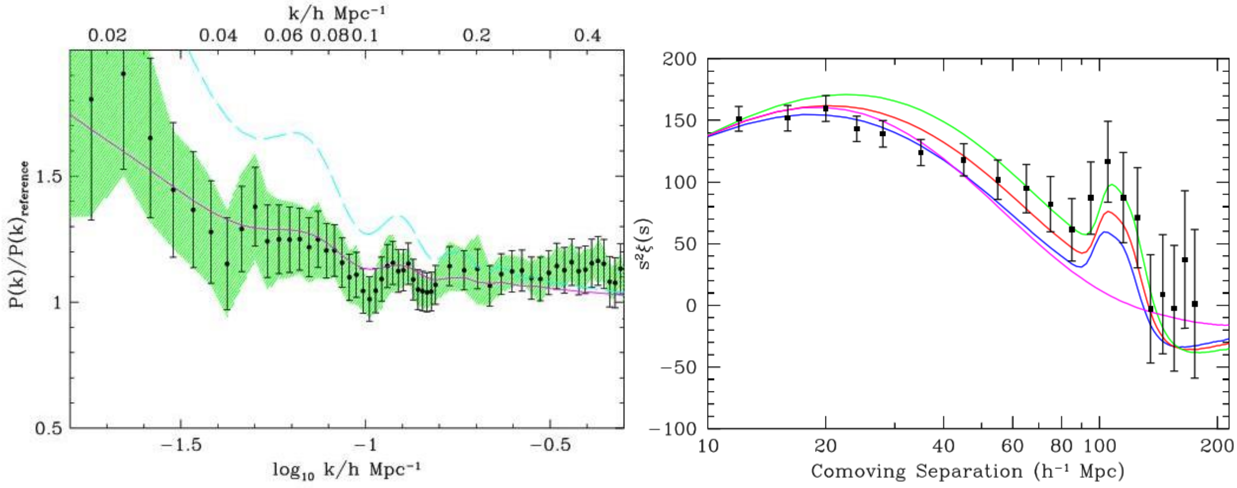

Figure 15: The detection of BAO in the 2dF and SDSS galaxy redshift surveys. Left: the power spectrum from the 2dFGRS, showing the BAO as the subtle wiggles in the curve. The wiggles correspond to the constructive and destructive interference of the behaviour of the sound waves with the underlying density perturbations from inflation. The magenta curve shows the prediction from theory; the blue dashed curve is the theory prior to corrections for the size of the 2dF survey. Right: the correlation function from the SDSS, showing the BAO as the acoustic peak at 100 Mpc/h (500 million light years). The green, red, and blue curves show the predictions from theory for three mildly different cosmologies. The magenta line shows a cosmological model without any BAO; one can see that the data favour the models with an acoustic peak (Taken from Cole et al. 2005 and Eisenstein et al. 2005).

BAO and the distance scale

Most important, these detections opened the door for the use of the BAO to measure the cosmic distance scale. Because the acoustic scale of 500 million light years can be predicted with accuracy, we can compare that scale to the one measured in our analyses based on the angular and redshift separations of galaxies. However, the latter measurement depends on the distance scale assumed in the analysis: if we change the cosmological model to alter the relation of redshift to distance, then the inferred separations of the galaxies will change. For example, if we double the distances to all of the galaxies, all of the galaxy separations will double. Hence, if the measured acoustic scale doesn't match to the predicted one, we can infer that our assumed distances are wrong. Adjusting until they match gives a measurement of the cosmic distance scale. This is an example of what is known as a standard ruler method.

The initial detections from SDSS and 2dFGRS provided a 4% measurement of the position of the acoustic peak and thereby a first measurement of the cosmic distance scale from BAO. This combination of the BAO data with the CMB data provided a strong argument for a spatially flat Universe with a matter density about 25-30% of the critical density − and therefore a dominant 70-75% contribution from dark energy.

Since these first detections, the study of the BAO has raced forward, with additional analyses of larger SDSS data sets and new detections from other surveys. The largest survey to date is that of the SDSS-III Baryon Oscillation Spectroscopic Survey (BOSS). This survey includes 1.5 million galaxies reaching out to redshift of 0.7. In 2012, the analysis of the first 35% of the data set produced a 1.7% measurement of the distance to redshift 0.57. Recently in 2014, the analysis of 85% of the final data improved that result to 1.0%, the most precise extragalactic distance yet measured. Other BOSS data have produced a detection of the BAO in the clustering of the intergalactic medium at redshift of 2.3, yielding a 3% measurement of the distance scale at this much earlier epoch in the Universe.

Theoretical work has also progressed rapidly. Large cosmological simulations and perturbation theory focusing on the BAO have shown that the acoustic scale is largely insensitive to astrophysical complications from galaxy formation and late-time gravitational structure formation. This robustness arises from the very large size of the acoustic scale: 500 million light years is big even in the context of galaxies and galaxy clusters. This is a great advantage for the method: we expect systematic errors to be well below 1% and indeed likely to be subdominant even for the yet more ambitious galaxy surveys of the coming decade.

The combination of the BOSS data and the latest CMB anisotropy data from the Planck satellite provides a compelling measurement of the cosmological model, based dominantly on the utilization of these sound waves from the early Universe. The data strongly support the inflationary prediction of a spatially flat Universe as well as the simple model of the dark energy as a cosmological constant. Combining with recent Type Ia supernova data, we reach measurements of the cosmic density (Ω_m) and the Hubble constant (H_0) with uncertainties around 2%.

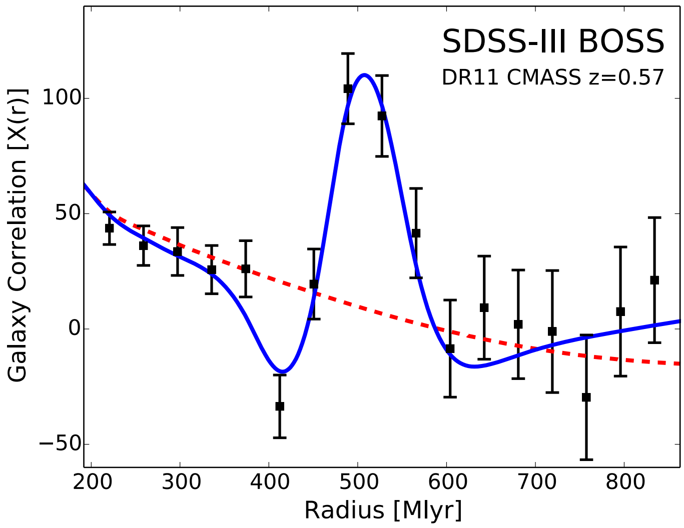

Figure 16: The latest correlation function from SDSS-III BOSS, showing a strong detection of the acoustic peak. In this figure, the correlation function statistic has been adjusted so that all of the error bars are independent from each other. The red dashed line shows the model without BAO; the blue line is the best-fit BAO model. The resulting statistical significance of the acoustic peak is about 8σ, an unmistakable detection. These data result in a 1.0% measurement of the distance to redshift 0.57 and provide a critical measurement of our cosmological model. Figure from Anderson et al. (2014).

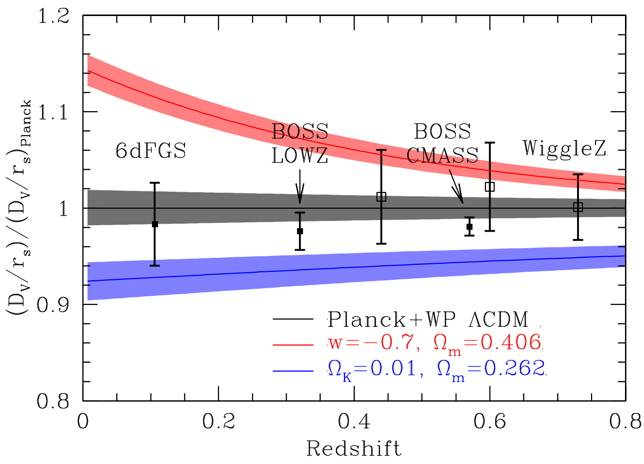

Figure 17:. Measurements of the cosmological distance scale from the BAO method from a variety of recent surveys. These distances have been divided by the spatially flat cosmological model with cold dark matter and a cosmological constant that best fits the data from the Planck satellite measurements of the cosmic microwave background anisotropies. The grey swath shows the variations allowed by Planck data (at 68% confidence). The red and blue swaths show how the distance predictions from Planck would change if one instead used a model with evolving dark energy or a slight spatial curvature. One can see that the BAO measurements have excellent ability to distinguish these models. With a full analysis, we find that the combination of Planck data and all BAO data produces results that are beautifully consistent with the simple model of a flat Universe with a cosmological constant. Figure from Anderson et al. (2014).

5 The future of Large-Scale Structure

The history of large-scale structure as related above is a great scientific success story, and has helped guide us to the present standard model of cosmology: a universe of critical density, divided between dark energy, dark matter and baryons in the rough proportions 70:25:5. Where do we go from here? How much interest is there in refining these numbers further? The first thing to say is that the task of surveying the universe is set to continue, as shown in Figure 18. This reveals a kind of observational Moore's Law, in which the number of measured redshifts doubles every 3-4 years. At present, we know about five million redshifts, but the total is expected to reach perhaps 100 million by 2025, based on projects that are underway. The largest of these are the ground-based DESI (Dark Energy Spectroscopic Instrument), which will use the Kitt Peak 4m telescope with a 2dF-like positioner for about 5000 fibres, and the European Space Agency's Euclid satellite, due for launch in 2020.

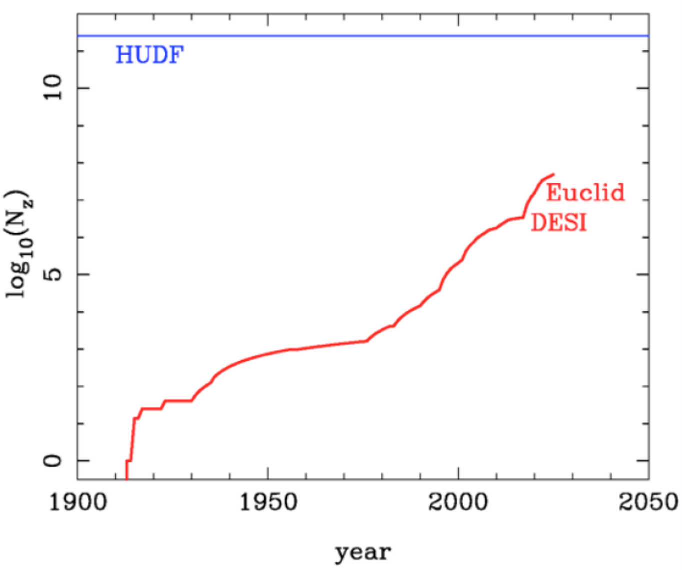

Figure 18: The growth in the number of galaxy redshifts with time, with an extrapolation to future projects. The line labelled HUDF indicates the ultimate limit in which we know the distance to every galaxy in the visible universe (to the depth of the Hubble Ultra Deep Field).

The scientific goals for these massive projects are motivated above all by the existence of dark energy, which provokes two central questions:

(1) Is dark energy simply an energy density of the vacuum? In that case it should be completely unchanging with time. Larger redshift surveys should allow us to measure the dark energy density at a number of different times, in order to test this prediction.
(2) Is dark energy an illusion? We infer its existence as a physical substance by evidence from BAO and other probes that the expansion of the universe is accelerating, plus assuming the correctness of Einstein's relativistic theory of gravity. It is therefore a logical alternative that a different theory of gravity might yield acceleration without dark energy.

Evolution of dark energy

The search for any change in the level of dark energy has been a major target in recent cosmological research. What we need is the change of the expansion rate with redshift, *H(z)*, the square of which is directly proportional to the density of the universe, telling us how the dark energy density evolves. As we have seen earlier, the BAO standard ruler lets us infer *H(z)* by measuring the distance-redshift relation. The results are consistent with the expectation of vacuum energy: over the period when the universe has doubled its size, the dark energy density can have changed by no more than 5-10% (compared to a factor 8 change in the matter density). Future experiments, leading up to DESI and Euclid, will improve the sensitivity to the point where even a 1% change will be detected.

Testing gravity

If all we know is the accelerating expansion of the universe, it is not possible to tell whether dark energy is real, or whether gravity needs to be modified. But if gravity changes its properties on the scale of the whole cosmos, compared to what we measure locally in the solar system, then there must be a change with scale. One fruitful place to look for such a change is on the scales where we measure redshift-space distortions (10 Mpc or so, perhaps 0.1% of the size of the visible universe). At the time of the first 2dFGRS study, in 2001, this signature was seen as useful for measuring the density of the universe; but if we accept that this is now well pinned down from the CMB, then the amplitude of peculiar velocities can tell us if superclusters are collapsing at the rate expected in Einstein's theory. As with dark energy, there is so far no sign of a non-standard effect: the strength of gravity on these scales is within about 10% of the standard value.

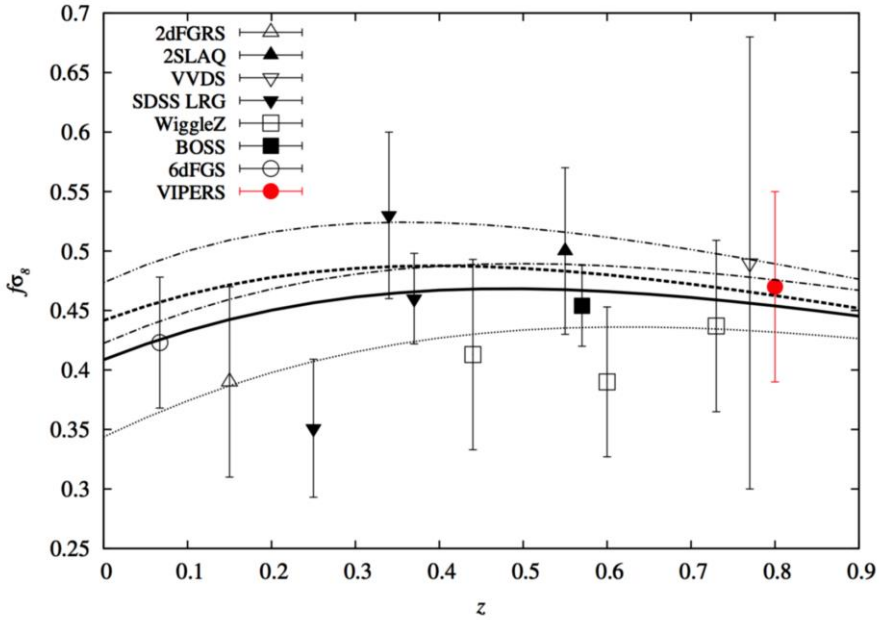

Figure 19: A plot of a quantity proportional to the rate of growth of density fluctuations with time, as a function of redshift. A large number of studies have extended the 2dFGRS redshift-space distortion measurements in precision and redshift coverage. All these results are consistent with the standard Einstein theory of gravity (solid line), and they are starting to rule out some of the altenative gravity models that are designed to explain the accelerating universe without invoking a cosmological constant. Data from Euclid will increase the precision to the 1% level, and extend the redshift coverage to *z*>2. Figure from de la Torre et al. (2013).

Outlook

It is possible that these studies will continue to yield null results even at the 1% level. Will we then be satisfied that we have a complete understanding of cosmology? Certainly not, as it is already clear that there are major puzzles connected with our standard model, particularly concerning the dark energy density. Simple calculations of the expected level of vacuum energy yield much larger numbers than observed − by between 60 and 120 powers of 10. The worry over this puzzle might be removed if we see the dark energy evolving: then it could have fallen to its present value from a much larger number. But making sense of a tiny static value is harder. One approach is to postulate a multiverse: perhaps everything we see is but one of many copies, each with a different vacuum density. But the formation of structure (down to galaxies, stars, planets and observers) can only happen if the vacuum density takes a small value: thus we would live in an unusual universe because only this permits life. Indeed, this argument was used in 1987 by Steven Weinberg to predict a small non-zero cosmological constant, at a time when most cosmologists were happy to assume that the vacuum was truly empty. The study of the large-scale structure of the universe therefore very much retains the potential for radical conclusions in cosmology.

Bibliography

Anderson, L., et al., 2014, MNRAS, 441, 24
Bennett, C. et al., 2013, ApJS, 208, 20
Beutler F. et al., 2011, MNRAS, 416, 3017
Blake C. et al., 2011, MNRAS, 415, 1707
Cole, S. et al., 2005,MNRAS, 362, 505
Colless, M. et al., 2001, MNRAS, 328, 1039
Davis, M. et al., 1982, ApJ, 253, 423
De la Torre, S. et al., 2013, A&A, 557, 54
Eisenstein, D.J. et al., 2005. ApJ, 633, 560
Eisenstein D.J. et al., 2011, AJ, 142, 72
Hubble, E. 1929, PNAS, 15, 168
Kaiser N. 1987, MNRAS, 227, 1
Maddox, S. et al. 1990, MNRAS, 243, 692
Peacock, J.A. et al. 2001. Nature, 410, 169
Peebles P.J.E. Yu J.T., 1970, ApJ, 162, 815
Sunyaev R.A., Zel'dovich Ya.B., 1970, ApSS, 7, 3
Shectman, S.A. et al., 1996, ApJ, 470,172
Weinberg, S., 1987, Phys. Rev. Lett., 59, 2607
York D.G., et al., 2000, AJ, 120, 1579