The Age of Things: Sticks Stones and the Universe
Ninth Lecture Notes
Matt Hedman
http://cfcp.uchicago.edu/~mmhedman/compton1.html
5/22/2004

Distances, Redshifts and the Age of the Universe

Based on a variety of evidence, the universe as we see it today does not appear to be infinitely old, but instead arose from a singular and formative event called the **big bang**, which happened at a definite, measurable time in the past. When the big bang occurred is one of the fundamental questions in cosmology (the study of the composition, structure and evolution of the universe on the largest observable scales) and recent observations have yielded important new constraints on the timing of this important event.

1 The dynamics of galaxies

Some of the most direct evidence for the timing and even the existence of a big bang comes from the distribution and dynamics of galaxies. **Galaxies** like our own Milky Way are more-or-less discrete clumps of stars and gas (and other forms of matter) which are scattered throughout the universe. Since these objects can be seen from billions of light years away, studying their locations and movements of the objects provides important information about the large-scale dynamics of the universe.

1.1 Redshifts

The redshift of a galaxy is a parameter that can be derived from an examination of the galaxy's spectra. For example, consider the spectra illustrated in Figure 1, which shows a number of sharp spikes at well-defined wavelengths. These features arise due to the presence of specific atoms and molecules in the galaxy. The quantum mechanical structure of an atom causes it to absorb or emit light very strongly at certain discrete wavelengths and produce narrow features in the spectra like those observed in the figure. Each element produces a distinct pattern of spikes that allows the features in the spectra of the galaxy to be associated with the presence of a particular element.

To match the features in the spectrum of this galaxy with those found in laboratory measurements of atomic gases, we must multiply the wavelength of the features found in the laboratory by 23%. All the features in the galaxy spectra are shifted by this factor, so it is unlikely that we have simply misidentified the elements responsible for these spikes. Instead, it appears that all the light from the galaxy has a 23% longer wavelength when it is measured here on earth compared to what it had when it was emitted from the atoms in the galaxy (where the atoms would have produced the same features we observe in laboratories on earth).

This 23% change in the wavelength of the light is reffered to as the **redshift** of the galaxy, so named because it is usually found that the features in the galaxy are shifted to longer wavelengths, or the red end of the spectrum (Only a few very nearby

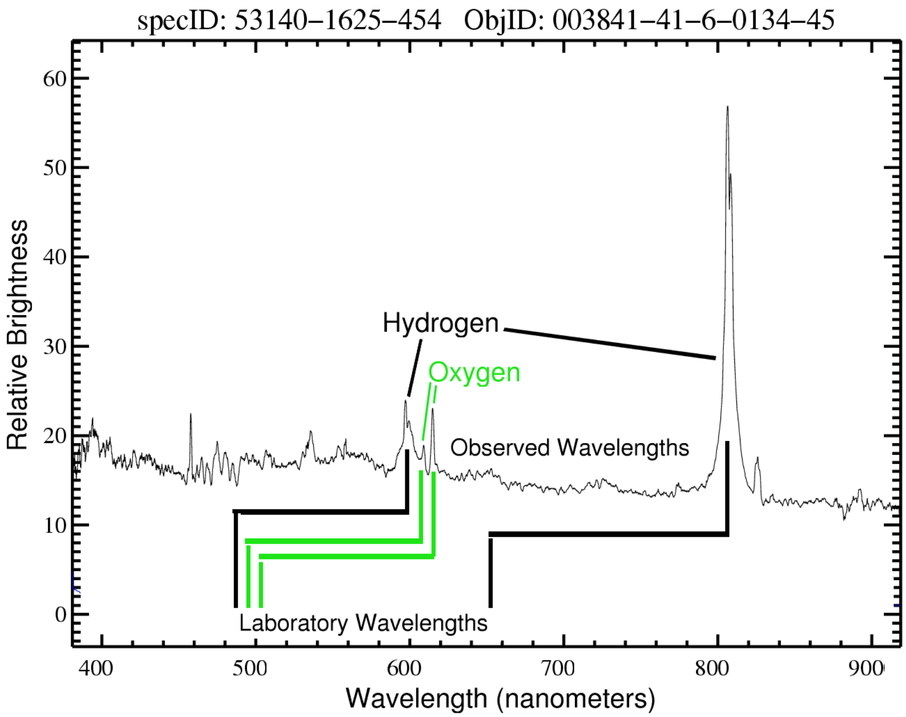

Figure 1: Spectra of a particular galaxy from the Sloan Digital Sky Survey (Courtesy of E. Sheldon). This plot shows the amount of light emitted by the galaxy at different wavelengths.These data show a number of narrow spikes that correspond to discrete transitions in various atoms. Spikes due to hydrogen and oxygen are identified and marked based on the unique spacing of these features. The observed wavelengths of these features are all shifted to longer wavelengths compared to the wavelengths of such features observed in the laboratory (indicated by the lines near the bottom of the plot).

galaxies like Andromeda have blueshifts). In standard astronomical notation, redshift is denoted by the letter z, so in this galaxy $z = 0.23$.

As we will see below, there are several possible way that light from a galaxy can acquire a redshift. For the moment, however, we can interpret the redshift as due to relative motion between the galaxy and us. A similar effect is commonly found in sound waves, such as a train whistle or a car siren, which have a noticeably higher pitch (corresponding to a shorter wavelength) when the vehicle is approaching compared to when it has passed by. If we interpret the redshift of the above galaxy as due to its movement, then the fact that the wavelength of the light has been shifted to longer wavelengths implies the galaxy is moving away from us (or, equivalently, we are moving away from it). The speed at which the galaxy is moving away from us can also be computed from the magnitude of the redshift. Since nearly every galaxy with a well-studied spectrum has a noticeable redshift, nearly every galaxy appears to be moving away from us at great speeds (from hundreds of kilometers per second on up). This result is somewhat curious, and becomes even more interesting when the distance to the galaxies is taken into account.

1.2 Measuring Distance

While it is relatively straightforward to measure the redshift of a galaxy and then infer how fast it could be moving away from us, determining the distance to the galaxy is a rather more challenging task. Parallax, discussed in the last lecture, provides a relatively direct measure of distance based on simple trigonometry. Unfortunately, this method cannot be applied to objects outside our own galaxy (the apparent motion of the objects is too small to be measured if they are more than about 1000 light years away). Other techniques are therefore required to estimate the vast distances separating galaxies.

As noted in the previous lecture, the apparent brightness of an object depends on how much light the object generates (its **luminosity**) and how far away the object is (if an object is located 100 meters away, it appears 4 times brighter than it would if it were 200 meters away). Therefore, if we know how much light the object generates we can use its apparent brightness to calculate how far away it is. The challenge is to find an astronomical object that not only is bright enough to be seen in other galaxies, but also has a known or calculable luminosity.

In practice, we cannot safely calculate the luminosity of any astronomical object from first principles. Fortunately, there are certain types of objects where the luminosity can be estimated based on other characteristics of the observed light. A good example of such a class of objects are the **Cepheids**, a type of star whose luminosity varies with time in a characteristic cycle. Cepheids first increase in brightness relatively quickly, then their brightness more slowly drops back down to its original level before the cycle repeats itself. The period of this cycle can range from days to weeks. These stars are, on average, thousands of times more luminous than the sun, so they can be observed in other galaxies. Indeed, hundreds of them have been found in the Magellanic Clouds, small satellite galaxies near to our own Milky Way. Since all the Cepheids in either of these clouds are about the same distance away, any variation in their brightness should due only to the differences in their luminosity.

The brightness of the Cepheids in the Large Magellanic Clouds vary over about 4 magnitudes, or a factor of about 40 (see the last lecture for a discussion of magnitudes). Clearly, the luminosities of all Cepheids are not the same. Fortunately, as shown in figure 2 the brightness of a Cepheid is very tightly correlated with the period of time it takes for it to undergo one cycle in brightness. This means we can use this period to infer how luminous the Cepheid is. For example, say we find a Cepheid in another galaxy with a period of about 10 days and a mean magnitude of 25. This Cepheid it then ten magnitudes, or 10,000 times fainter than a Cepheid with a comparable period in the Large Magellanic Cloud. The Cepheid, and the galaxy that contains it therefore must be 100 times farther away than the Large Magellanic Cloud (remember, brightness drops a factor of four every time the distance increases by a factor of two).

If we want to know how far away the galaxy is in light years, then we just have to know the distance to the Large Magellanic Cloud. Fortunately, there are Cepheids in our own galaxy which are close enough that we can measure their distance using the parallax method. This allows us to determine the luminosity of these Cepheids and infer the luminosity of the Cepheids with the same periods in the Magellanic Cloud. This sort of analysis indicates the Large Magellanic Cloud is about 150 thousand light

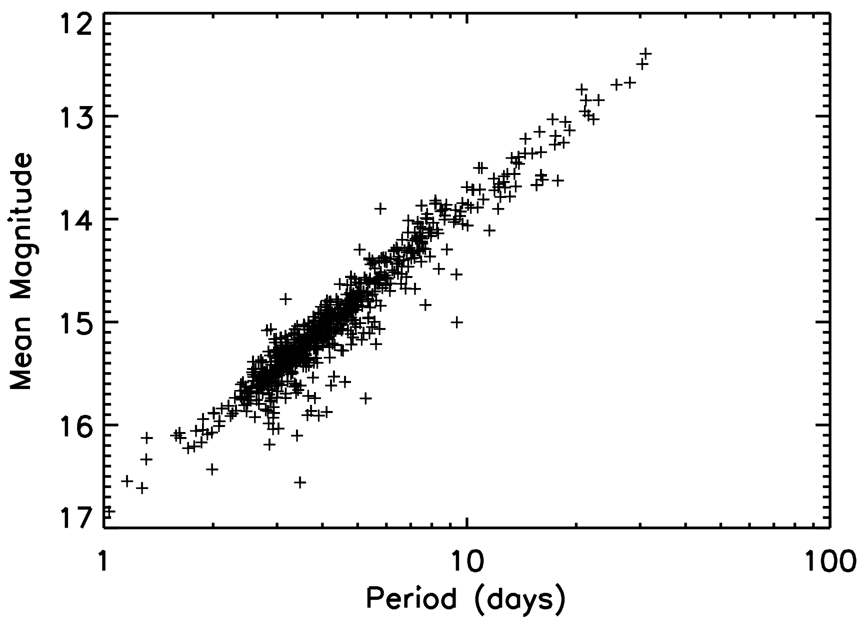

Figure 2: The relationship between the mean brightness and the period of brightness variations for Cepheids in the Large Magellanic Cloud (based on data from Udalski et. al. in *Acta Astronomica* Vol 49 (1999) pg 223). Since all of these Cepheids are the same distance away from us, the variations in brightness reflect variations in the luminosity of the Cepheids. The brightness of these objects is correlated with the period, so by measuring how fast a Cepheid changes its brightness, we can get a reasonable estimate of its luminosity.

years away (and other methods of measuring the distance to this object have yielded roughly similar results). Since the galaxy in the above example is 100 times further away than this, we can therefore deduce it is about 15 million light years away.

While Cepheids are quite bright compared to the sun, they still can only be identified in galaxies within about 100 million light years of us. Many galaxies exist beyond this, and much brighter objects are needed to estimate the distances to these objects. Recently **Type Ia supernovae** have emerged as powerful tools for measuring such great distances.

Supernovae produce huge amounts of light for a short period of time. For about a week, they can be as luminous as a billion suns. Type Ia supernova are a specific type of event which can be identified by particular features in their spectra. (These events are believed to occur because some star dumps material on a white dwarf, which restarts the nuclear reactions in the dwarf and causes it to explode). While supernovae are relatively rare (occurring about once a century in any galaxy), a number of these events have now been observed in galaxies at known distances based on Cepheids and other data. This means we can estimate the total amount of light generated by the different Type Ia Supernovae. These measurements show that there is some variation in the peak brightness of these events, but that this variation is correlated with how quickly the supernova brightens and dims. Thus, as with the Cepheids, we can infer the peak luminosity of the Supernova by measuring how its brightness changes with time. Given this luminosity, the apparent brightness of the supernova allows us to

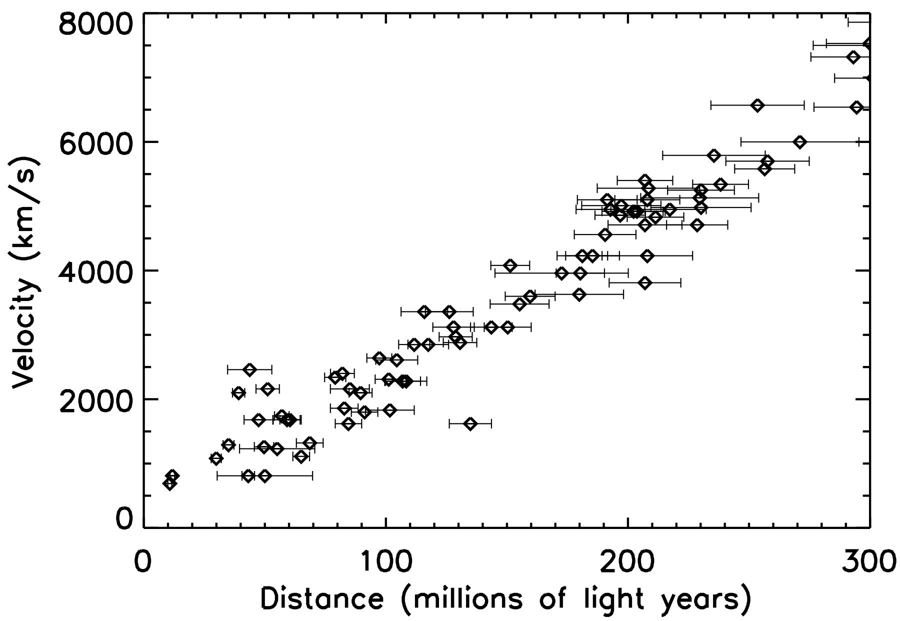

Figure 3: A Hubble Diagram showing the velocity inferred from the redshift of the galaxy versus the distance deduced the Type Ia Supernova brightness (based on data from Tonry et. al. 2003). The horizontal error bars show the uncertainty in the distance measurement in each cluster (uncertainties in the redshift are much smaller). Note that more distant galaxies are moving away from us faster than more nearby galaxies

compute the distance, which can be billions of light years.

2 Hubble Diagrams

Figure 3 shows how fast some galaxies appear to be moving away from us (based on the redshift) as a function of their distance away from us (based of Supernova brightness measurements). This **Hubble Diagram** shows that the apparent velocity of a galaxy is directly proportional to its distance away from us. This situation can be illustrated as follows:

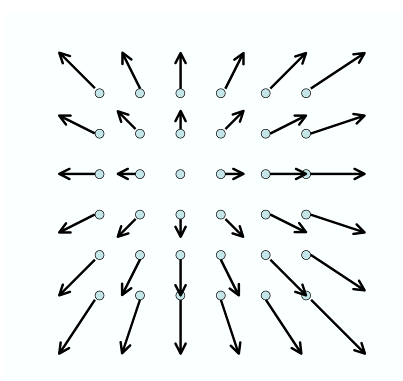

The gray dots here represent galaxies. The galaxy with no arrow is the one we are living on and the arrows on the other galaxies indicate how quickly and in what direction they appear to be moving away from us.

Now we do not imagine that we are so cosmically unpopular that all galaxies are moving any from us in particular, and indeed we do not this to be the case to explain the data in the Hubble diagram. For example, imagine all the galaxies were flying away from some single point, with a speed proportional to their distance from that point, as shown here:

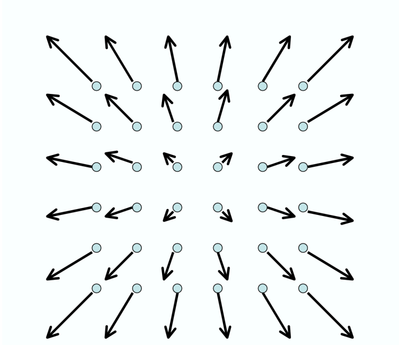

In this case, any galaxy in this group would observe the same relationship between distance and relative velocity.

As we will see in the next section, this picture of galaxies flying through space from a single point is probably not the correct way to visualize how the universe behaves. Even so, this naive picture is enough to suggest that some singular event happened did happen deep in the past. The Hubble diagram shows a galaxy 100 million light years away is moving away from us at about 2000 kilometers per second. If we assume this galaxy has always moved at roughly the same speed at which it is moving now, then we can calculate how far away this galaxy was at away from us at any other time in the past. Since the galaxy is moving away from us, in the past it was closer to us, and it works out the 15 billion years ago there was no separation between us and this galaxy. Now consider a galaxy twice as far away (200 million light years). It is travels twice as fast (4000 kilometers per second), so 15 billion years it will be be right on top of us as well. Indeed, if we assume all the galaxies moved at a constant speed, all the galaxies would have been essentially in the same place 15 billion years ago. It is therefore very tempting to imagine that galaxies were thrown out from some gigantic explosion somewhere is space 15 billion years ago. However, this picture is almost certainly wrong.

While there is nothing in the data from the Hubble diagram which explicitly contradicts this idea of galaxies moving through space from some central point, such a model does not mesh with other observations of the large-scale structure of the universe. If all the galaxies were launched from a special and specific point in space, then we expect that the characteristics and distribution of the galaxies would depend on their distance from this central point. For example, we might expect that galaxies further from this point would be lower mass than galaxies closer in (or be a different type, or have a different spectra), or there would be more galaxies nearer to this point than further away. Thus far, no observation of the large-scale structure of the universe shows any pattern like this. Indeed, the universe on large scales appears to be basically homogeneous, with the same types of galaxies, in roughly similar distributions, occurring throughout space. (There are definitely clumps and groups of galaxies of a variety of sizes in various places, but these structures do not extend to fill the entire observable universe.) This strongly contradicts the notion that there could some special point in the universe where all the matter could come from.

The large-scale homogeneity in the structure of the universe indicates that the patterns in the apparent motions of galaxies are due to a more broad-scale phenomenon, which affects all regions of space equally. Just this sort of phenomenon occurs naturally in the context of general relativity.

3 General Relativity and the Expanding Universe

General relativity is the theory which currently provides the most accurate method for calculating how objects move under the influence of gravity. While the mathematical manipulations required to extract predictions from this theory are quite intimidating, the basic premise behind these computations is straightforward, although it might seem strange: Gravity is not a force so much as a distortion in the geometry of space and time.

In the classical mechanics of Newton and Galileo, an object moves along at a constant speed in a straight line unless acted upon by an outside force. An outside force causes an acceleration (a change in the speed or direction of the object's movement) which depends on the mass of the object. In most situations, the force couples differently the different objects, so different objects will move along different paths (For example, an electric field will deflect a positively charged proton one way, deflect a negatively charged electron another way, and not deflect a neutron at all). Therefore in general we need to know something about the characteristics of the object (such as its charge and mass) to compute how it will move in a certain situation.

Gravity is anomalous in that we do not need to know anything about the composition or mass of the object to calculate how it will respond to a gravitational field. A famous demonstration of this involves a feather and a lead ball. Both objects are held above the surface of the earth in an evacuated chamber. When these objects are released, they fall at the same rate.

In classical physics, this result is something of a coincidence. The mass of an object determines both the strength of the gravitational force and the how quickly the object accelerates in response to the force. These two effects cancel each other out in the equations that describe how the object actually moves.

In general relativity, on the other hand, the fact that the characteristics of the object are irrelevant to the calculation of its movement is fundamentally important. It means that the motion of bodies in a gravitational field has something in common with the motion of bodies when no outside force is operating. When there are no outside forces acting on a particle, it moves at a constant speed in a straight line, regardless of its mass or any other characteristics it may have. A straight line is of course the shortest distance between two points, so the path the particles take can be deduced based only on geometry. General relativity posits that massive objects distort the geometry of space (and time) and alters what is the effectively the shortest distance between any two points. In the absence of other forces, all objects continue to follow this new, geometrically defined path, but because of the distortion, they may not appear to follow a straight line. General relativity thus suggests that gravity is not a force which causes particles to deviate from a straight-line path, it is a change in the definition of what a straight-line path is.

(As a side note, I should point out that the relevant displacement in general relativity actually involves changes in both space and time. If we took a ruler and measured the number of kilometers traveled by a spacecraft as it traveled in a curved orbit around earth, we would *not* find that it traveled a smaller number of kilometers than a spacecraft which traveled along a straight line between the same points. However, if we also took into account the amount of time elapsed on a clock on the spacecraft in orbit and the spacecraft traveling along a straight line, the total displacement in time and space would be at an extreme for the spacecraft in orbit)

General relativity is not just an intriguing new way to look at gravity, it also explains phenomena which could not be understood using classical gravity (such as irregularities in the orbit of Mercury) and predicted effects (such as the bending of starlight by the sun and other massive objects) that were later confirmed by observations. This theory therefore provides the best basis we have for understanding how gravity operates.

General relativity also yields an explanation for how we can obtain a Hubble diagram like the one we observe in a homogeneous universe. Imagine a universe filled with a homogeneous distribution of matter, such as a reasonably evenly distributed set of galaxies, as soon here (of course, imagine the pattern continuing forever in every direction):

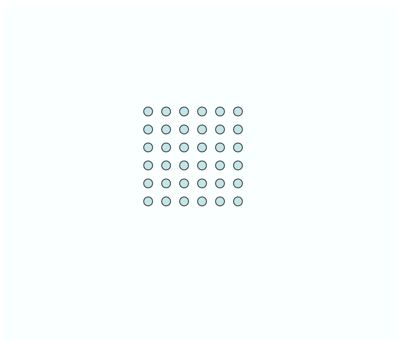

As this matter alters the geometry of the universe, it changes the effective distance between particles and the density of the material, which further alters the geometry of the universe. The distances between any pair of galaxies is therefore dynamic and can expand or contract like a rubber band. For example, after some amount of time the spacing between the galaxies could have increased to this:

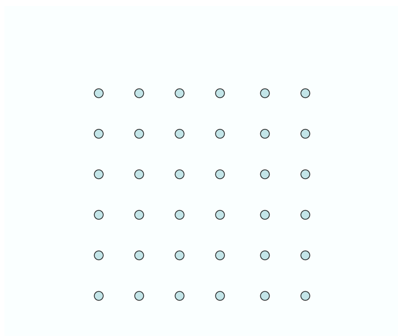

This change in the distances between the galaxies has some similarities with the explosive model discussed above. However, unlike this model, the galaxies are not moving through space, but the amount of "space" between the galaxies is increasing.

Since this expansion is happening throughout the universe, there does not need to be a special central location where all the material originally came from. The universe can therefore remain more-or-less homogeneous, as we observe.

3.1 Re-interpreting the Hubble Diagram

If the universe is expanding as we expect in the context of general relativity, then concepts like distance and velocity are not so easy to define. Therefore the connection between the observable parameters in the Hubble diagram and the dynamics of the universe must be re-examined.

For one, it is no longer so useful to interpret a redshift in terms of a relative velocity between us and the galaxy, since galaxies are not exactly moving through space in this scenario. Instead, changes to the wavelength of the light occur because as this light propagates through the universe, its wavelength expands or contracts along with the universe (Note that when I say the universe is expanding, I mean the distance between any given pair of galaxies is increasing). If the universe expands by 20% in the time it takes the light to travel from some galaxy to us, then the wavelength of the light will be 20% longer when we measure it here on earth compared to what it was when it was generated by the atoms in the galaxy and the galaxy then has an observed redshift of 0.20.

It therefore makes sense to re-draw the Hubble Diagram with the actual redshift along the y-axis instead of an apparent "velocity" (also, we now rename the distance the luminosity distance, anticipating issues discussed below):

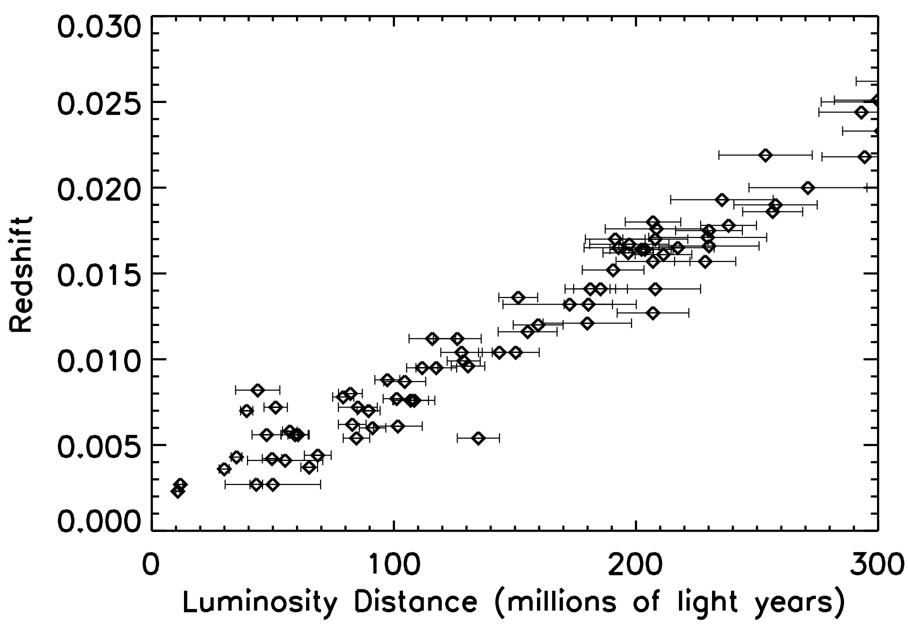

This plot shows that a galaxy 100 million light years away has a redshift of 0.01, so the universe had expanded by 1% in the time it took the light to travel from that galaxy to us. The spacing between any pair of galaxies was therefore 99% of its current value when the light left that galaxy. This change is the amount of space between galaxies is usually quantified using a **scale factor**, which is the distance between any two galaxies at some time, divided by the same distance today. For the above galaxy, the scale factor would be 0.99 when the light was emitted by the galaxy.

Since the scale factor of the universe when the light was emitted from the galaxy can be computed directly from the redshift of the galaxy, we can easily make the following plot:

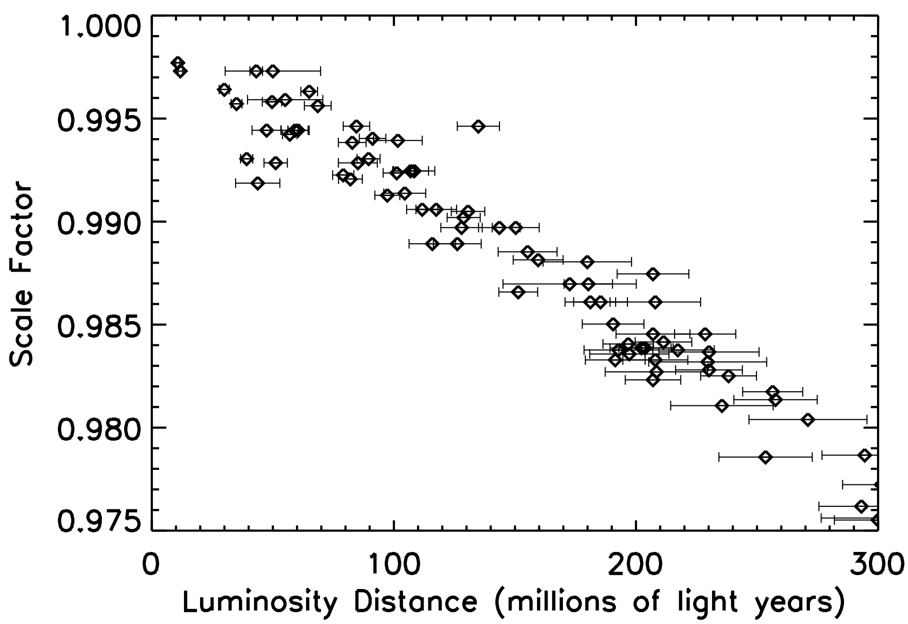

This plot shows how the scale factor of the universe when the galaxy emitted the light varies as a function of the *distance* to the galaxy. However, what we would really like is a plot showing how the scale factor has varied as a function of *time*. Since light travels at a finite speed, it should take light a longer time to reach us from more distant galaxies. Since a light year is the distance light travels in a year, if the universe was not expanding, we could simply say that the galaxy 100 million light years away emitted its light 100 million years ago.

Of course, the universe is expanding, which complicates things enormously. As the light travels from the galaxy to us, the distance between the galaxy is constantly increasing, so the distance the determined by the luminosity of the galaxy (the so-called **Luminosity Distance**) used here is not simply the speed of light times the time it took the for the light to make the trip. It also depends on how the universe has expanded (that is, how much the scale factor has changed) over this time.

The expansion of the universe has only a small affect for the data shown in the plot above because the scale factor does not change very much. However, recently Supernova have been observed in galaxies with redshifts exceeding one. The universe has therefore expanded by over a factor of two during the time between when this light was emitted by the galaxy and when it was observed on earth, and the effects of the changing scale factor are very important. If we plot the scale factor versus luminosity distance for these new observations (compiled in Riess et. al. 2004), we get this:

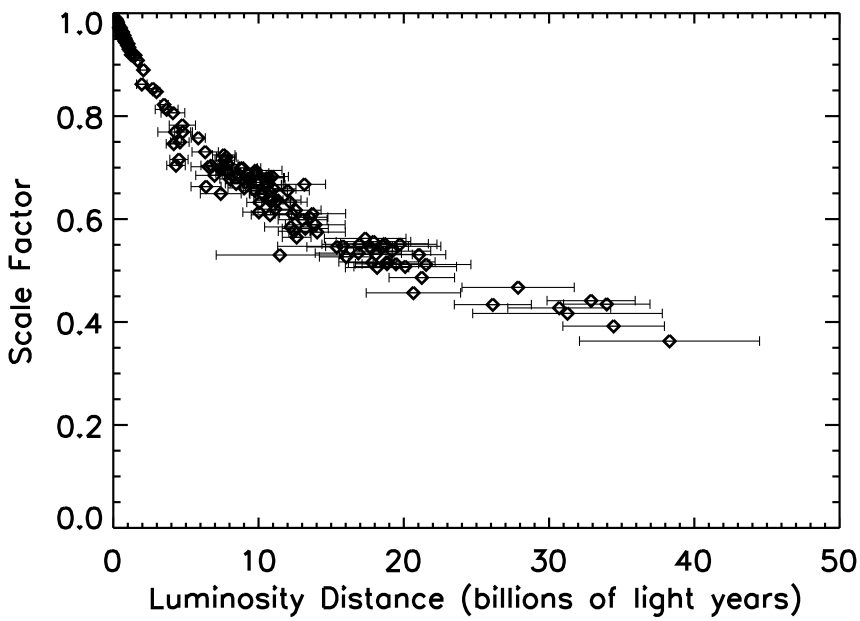

Where the "distance" of the furthest galaxies is 30 billion light years. However, these apparent distance are heavily inflated due to the large changes in the scale factor. We can account for this and deduce the time it really took the light to travel between the galaxy and us (making a few assumptions and approximations):

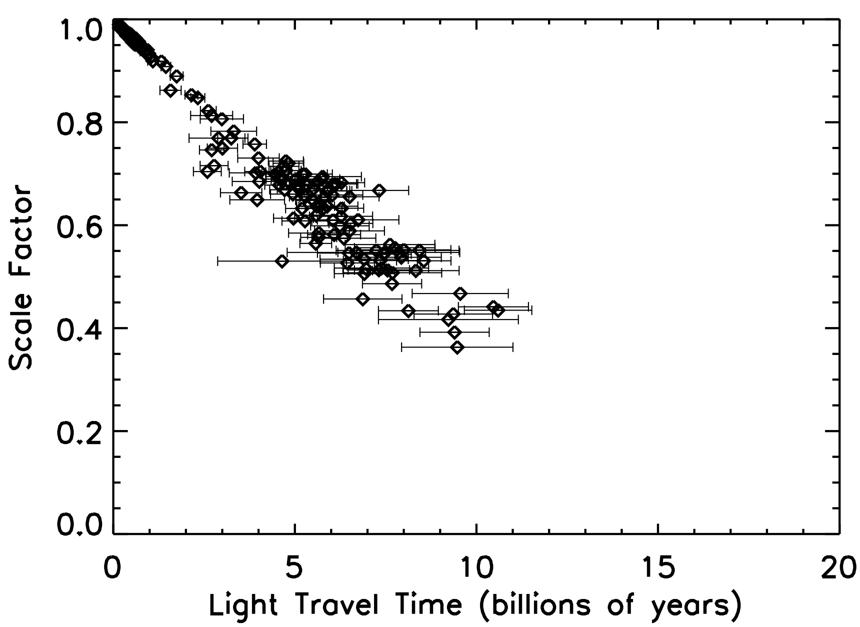

This plot then finally gives the scale factor as a function of the time when the light was emitted. Clearly, as we go further into the past, the scale factor gets progressively smaller. Indeed, the data are more-or-less following a straight line, showing the scale factor has increased with time in a more-or-less constant rate. Extrapolating this trend further back in time, we find the scale factor would be zero about 15 billion years ago. A scale factor of zero corresponds to zero distance between adjacent galaxies, so the entire universe was infinitely dense at this time. This singularity in the size and density of the universe is the real "big bang" that marks the formative moment in the history of the universe. (It is amusing to note after all this work we get basically the same estimate for the age of the universe as we did with the overly simplistic considerations above).

While the timing of the big bang provides a reasonable basis for measuring the age of the universe, the age obtained by fitting a straight line to the data and extrapolating all the way to zero is not entirely accurate. Indeed, the available data do not exactly follow a straight line, but instead curve slightly. This is to be expected, since the amount and type of matter in the universe affects how quickly the universe expands with time. Therefore, to obtain an accurate estimate of the age of the universe, we need more data on the composition of the universe. Fortunately, as we will see in the next lecture, a variety of cosmological observations now provide important constraints on the matter and energy content of the universe and yield very precise estimates of the age of the universe.

4 References

For a general introduction to cosmology and the expansion of the universe, see the relevant chapters in:

- Freedman and Kaufmann *Universe, Sixth Ed* (Freeman and Co, 2001)

There is also a Special Edition of *Scientific American* entitled "The Once and Future Cosmos" that may be worth reading. Also the February 2004 issue of *Scientific American* has an article on the new Supernova results.

For the intrepid reader interested in learning how to solve problems in General relativity, try:

- Hartle *Gravity, An Introduction to General Relativity* (Addison-Wesley 2003)
- B. Schutz *A First Course in General Relativity* (Cambridge U Press 1994)

For measuring distances with Cepheids, see

- W.L. Freedman et al "Final results from the *Hubble Space Telescope* Key Project to Measure the Hubble Constant" in the *Astrophysical Journal* 553:47-72

For the latest on Type Ia supernova measurements, see:

- R.A. Knop et. al. "New Constraints on...." in *Astrophysical Journal* vol 598 (2003) pp 102-137.
- A.G. Riess et al "Type Ia Supernova discoveries at z>1...." `astro-ph/0402512` available at `www.arxiv.org`