8.942 Cosmology Notes

Andy Lutomirski

This copy is from 2008-11-18 15:32:35 -0500 (Tue, 18 Nov 2008).

Contents

1 Introduction

These are my notes from Max Tegmark's 8.942 (Cosmology) lectures at MIT in the fall of 2008-09. I'm typing them in class and they aren't very well edited, so they almost certainly contain errors and omissions. I hope that they are useful nonetheless.

If you want a copy of the LyX or LaTeX source, ask me.

1.1 GR

The formula for gravitational redshift is

$$\frac{\Delta\nu}{\nu} = \frac{gh}{c^2} = \frac{\phi}{c^2} \to \frac{d\tau}{dt} = (1+\phi)\,,$$

so the metric for Newtonian gravity is

$$d\tau^2 = (1+2\phi)\,dt^2 - dx^2 - dy^2 - dz^2.$$

Objects in free fall follow geodesics — that is, they try to maximize their own aging. So a falling object has to balance time dilation with the gravitational redshift. The total aging is

$$\begin{aligned}
\tau &= \int d\tau = \int_A^B \sqrt{(1+2\phi) - \dot{x}^2 - \dot{y}^2 - \dot{z}^2}dt \\
&= \int_{t_A}^{t_B} f(t, x, \dot{x})\, dt \\
&\approx \int_{t_A}^{t_B} \left(1 + \phi - \frac{\dot{x}^2}{2}\right) dt.
\end{aligned}$$

2 Zero-order cosmology

2.1 The FLRW metric

The FLRW metric is

$$\begin{aligned}
d\tau^2 &= dt^2 - a(t)^2 \left(\frac{dr^2}{1-kr^2} + r^2 d\theta^2 + r^2 \sin^2\theta d\phi^2\right). \\
&= dt^2 - a(t)^2 \left[\frac{dr^2}{1-kr^2} + d\Omega^2\right],
\end{aligned}$$

where $k = 0, \pm 1$ and, if $k = 1$, r is bounded.

Hubble noticed that objects seemed to recede at velocity proportional to their distance. (There were all kinds of problems with his paper, but he was right in his conclusion. Einstein objected to Lemetre's idea of an expanding universe at first—there was a tendency for people not to believe their own equations.)

An object is said to be comoving if its (r, θ, ϕ) coordinates are constant. The distance scales as $d = d_0 \frac{a(t)}{a_0(t)}$, so $v = \dot{d} = d_0 \frac{\dot{a}(t)}{a(t_0)}$ and $\frac{v}{d} = \frac{\dot{a}}{a} = H$, the Hubble parameter.

$$H = \frac{d}{dt} \ln a \tag{1}$$

H used to be called the Hubble constant, but it isn't constant. The other major parameter is the redshift $z = \frac{a_0}{a} - 1$.

If you plug the FLRW metric into the EFE, you end up with the Friedman equation

$$H^2 = \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{kc^2}{a^2} = \frac{8\pi G}{3}\rho_{\text{tot}}.$$

The solution depends on what kind of stuff there is. We'll consider matter, radiation, curvature, and the vacuum.

Matter $\rho \propto \frac{1}{a^3}$ to conserve mass. This works for normal matter and for dark matter.

Radiation $\rho \propto \frac{1}{a^4}$ because the number density falls off like a^{-3} but the redshift adds another factor of a^{-1}. This applies to any ultrarelativistic particles.

Vacuum ρ is constant. It turns out that this stuff has negative pressure, so energy is still conserved.

Curvature $\rho \propto a^{-2}$ by reading the equation. The sign of this term can vary.

We can solve the Friedman equation easily if we ignore all but one term. In the matter-dominated case, we get $a \propto t^{2/3}$. For radiation, $a \propto t^{1/2}$. For curvature, with $k = -1$, $a \propto t$, which turns out to be equivalent to Minkowski space. For vacuum-dominated space, $a \propto e^{Ht}$.

$H(z)$ is the most convenient parameterization, because z is easy to measure. We can write

$$H^2 = \frac{8\pi G}{3}\left(\rho_\gamma + \rho_m + \rho_k + \rho_\Lambda\right) = H_0^2\left[\Omega_\gamma (1+z)^4 + \Omega_m (1+z)^3 + \Omega_k (1+z)^2 + \Omega_\Lambda\right],$$

where $\rho_* \equiv \frac{3H_0^2}{8\pi G}$ and $\Omega_i \equiv \frac{\rho_i}{\rho_*}$ We can write $\Omega_{\text{tot}} = 1 - \Omega_k$.

To compute the ages of things, we write

$$\begin{aligned} dt &= \frac{d\ln a}{H} \\ \Delta t &= \int \frac{d\ln(1+z)}{H} \\ &= -\int \frac{dz}{(1+z)\,H(z)}. \end{aligned}$$

The cosmological principle is the idea that the universe looks the same from everywhere as it does from here. Max thinks its rubbish.

2.2 Derivation of the FRW metric

In 3D Euclidean space,

$$ds^2 = dx^2 + dy^2 + dz^2 = |d\bar{r}|^2 = d\bar{r}^T g d\bar{r},$$

and, in 4D Euclidean space,

$$ds^2 = dx^2 + dy^2 + dz^2 + d\mathring{a}^2 = |d\bar{r}|^2 + d\mathring{a}^2.$$

If we restrict ourselves to the top half of the surface, $|\bar{r}|^2 + {}^2 = 1 \implies \bar{r}\cdot d\bar{r} + \mathring{a}d\mathring{a}$, and

$$ds^2 = |d\bar{r}|^2 + \frac{(\bar{r}\cdot d\bar{r})^2}{1-\bar{r}^2}.$$

More generally, $ds^2 = |d\bar{r}|^2 + k\frac{(\bar{r}\cdot d\bar{r})^2}{1-k\bar{r}^2}$.

In polar coordinates, we write $\bar{r} = r\hat{r}$ and differentiate to get $d\bar{r} = \hat{r}dr + r\hat{\theta}d\theta + r\sin\theta\hat{\phi}d\phi$. After a certain amount of manipulation, we get

$$ds^2 = \frac{dr^2}{1-kr^2} + r^2 d\Omega^2$$

Robertson and Walker proved that this is the only isotropic and homogenous metric.

2.3 Conformal time

Light follows null geodesics, so $d\tau^2 = 0 = dt^2 - a(t)^2 \left[\frac{dr^2}{1-kr^2} + r^2 d\Omega^2\right]$. If we fire a photon radially, then

$$\begin{aligned} dt &= a(t) \frac{dr}{\sqrt{1-kr^2}} \\ \underbrace{\int \frac{dt}{a(t)}}_{\text{conformal time } \eta} &= \underbrace{\int \frac{dr}{\sqrt{1-kr^2}}}_{S_k(r)}. \end{aligned}$$

So we can rescale our coordinates by using conformal time and $S_k(r)$, light-cones go at 45° angles. As a result, the comoving time between light crests is the same at any two positions (along a radial path). This gives $\frac{\delta t_2}{\delta t_1} = \frac{a(t_2)}{a(t_1)}$. It turns out that this applies to the momentum of everything, relative to comoving observers: $\bar{p} \propto \frac{1}{a}$.

2.4 Testable implications

If you really believe in this stuff, there would be some interesting effects.

- Observed galaxy recession (Hubble's law)
- Big bang nucleosynthesis: when the universe was small and hot, fusion should have happened. This predicts the proportions of light elements correctly.
- The night sky is dark! The brightness of stars falls off as r^{-2} but the number goes as r^3, so in an infinitely old universe, the sky should be very bright.
- Distant objects look younger.

What is this evidence for, exactly? Max believes that our entire observable universe was once as hot as the core of the sun, doubling its size in under a second. But there is no evidence of a singularity or that anything was ever at the Planck density.

It's also a bit odd to say "*the* universe."

2.5 Controversies

Cosmology made two famous predictions. First, the age of the universe was about 2 billion years. Geologists had 3 billion year old rocks. Secondly, cosmology predicted the distribution of elements based on big bang nucleosynthesis and didn't predict heavy elements, which was "wrong" because astrophysicists knew that stars didn't get past carbon. Gamow predicted the microwave background and no one believed him for years, until Wilson and friends found it and thought it was noise.

2.6 Angular diameter distance and luminosity distance

Classically, brightness $\Phi = \frac{L}{4\pi d^2}$ and $\Delta\theta \approx \frac{D}{d}$. In the FRW metric, $d_L(z)$ and $d_A(z)$ change. In GR, the angular diameter distance is affected by curvature and by the change in size of the universe.

Suppose we look at something far away that subtends an angle $\Delta\theta$. We compute

$$
\begin{aligned}
d_\perp &= \int \sqrt{-d\tau^2} \\
&= \int a(t)\, r d\theta \\
&= a(t)\, r \Delta\theta \\
&= d_A \Delta\theta.
\end{aligned}
$$

This is useless because we don't know what $a(t)$ or r is for anything we look at.

We can apply the Friedman equation (using the convention $k = \pm 1$):

$$
\begin{aligned}
\rho &= \Omega_{\text{tot}}\rho_* = (1-\Omega_k)\,\rho_* = (1-\Omega_k)\frac{3H^2}{8\pi G} \\
H^2 &= H^2 - H^2\Omega_k - \frac{kc^2}{a^2} \\
\Omega_k &= \left(\frac{cH^{-1}}{a}\right)^2 k \\
&\approx -k\left(\frac{\text{radius of horizon}}{\text{radius of curvature}}\right)^2 \\
a_0 &= \pm\sqrt{-\frac{k}{\Omega_k}}cH^{-1} = cH^{-1}\left|\Omega_k\right|^{-1/2}
\end{aligned}
$$

We can then write $a(z) = \frac{a_0}{1+z}$. Using the definition of conformal time,

$$
\begin{aligned}
\Delta\eta &= \int \frac{dt}{a} = \int \frac{dt}{da}da \\
&= \left|\Omega_k\right|^{1/2} \underbrace{\int \frac{H_0}{H(z)}dz}_{\gamma} \\
\gamma &= \int \left[\Omega_\Lambda + \Omega_k(1+z)^2 + \Omega_m(1+z)^3 + \Omega_\gamma(1+z)^4\right]^{-1/2} dz \\
r &= S_k(\Delta\eta) \\
d_A &= \frac{a_0}{1+z}S_k(\Delta\eta) = \frac{cH_0^{-1}\left|\Omega_k\right|^{-1/2}}{1+z}S_k\left(\left|\Omega_k\right|^{1/2}\gamma\right) \\
&= \frac{cH_0^{-1}}{1+z}\frac{S_k\left(\left|\Omega_k\right|^{1/2}\gamma\right)}{\left|\Omega_k\right|^{1/2}}
\end{aligned}
$$

The incoming flux $\Phi = \frac{L}{4\pi d_L^2}$ where d_L is the luminosity distance. But in GR, we're affected by the redshift as well as the decrease in the rate at which photons arrive. The result is

$$d_L(z) = d_A(z)(1+z)^2 .$$

We can use the luminosity distance to probe the expansion history of the universe by looking at type 1A supernovae.

Using baryon acoustic oscillations, we determine that there's extra correlation of galaxy positions at a scale of 150Mpc. We can find these correlations and use them to compute d_A. We can do the same thing to the microwave background.

3 Inflation

3.1 The horizon and flatness problems

The air in the room we're in is homogenous, which is fine, because the air on one side of the room is in causal contact with the air on the other side of the room. But if we plot conformal time, we can see things now that were never in causal contact (if $t = 0$ is close to the CMB last scattering surface). For a ballpark estimate, the horizon is at $ct \sim cH^{-1} \sim 3GPc$. This is called the horizon problem.

If we have $\rho \sim a^{-3}$ or a^{-2}, then we expect curvature to dominate after awhile in the expanding universe. If $\rho \sim 1$, then the other stuff should dominate. Mathematically, $|\Omega_k| = \left(\frac{\dot{a}}{c}\right)^{-2}$. Depending on the content of the universe,

$$|\Omega_k| \propto \begin{cases} a^2 & \text{if radiation-dominated} \\ a & \text{if matter-dominated} \\ a^{-2} & \text{if vacuum-dominated} \end{cases} .$$

Suppose the universe were RD at the Planck time 10^{-43} seconds and became MD at $t \sim 10^{12}$ seconds. Then, during the RD phase, $|\Omega_k| \propto a^2 \propto t$ and grew by a factor of 10^{55}. Suring the matter dominated phase, $|\Omega_k| \propto a$ and grew by a factor of 3000 or so. We need $|\Omega_k| \lesssim 10^{-60}$ at the Planck time. This is called the flatness problem.

If the expansion history (ρ vs. a) crossed the line for curvature-domination that goes through the current state, then things that are in causal contact now were in causal contact at the beginning as well.

3.2 Inflation with a scalar field

We need to produce some kind of stuff that doesn't dilute. We begin with the Friedman equation and drop the curvature term because it becomes negligible quickly. (This is an attractor solution.) We define $\bar{m} = \frac{m_{\text{planck}}}{\sqrt{8\pi}}$, where $m_{\text{planck}} = \sqrt{\frac{\hbar c}{G}}$. $\hbar = 1$ for these purposes.

$$H^2 = \frac{8\pi G}{3}\rho - \frac{kc^2}{a^2} = \frac{\rho}{3\bar{m}^2}$$

Now we declare that

$$\ddot{\phi} + 3H\dot{\phi} + V'(\phi) = 0.$$

The field has units of mass. By looking at the Lagrangian (what Lagrangian?),

$$\begin{aligned} \rho &= \frac{1}{2}\dot{\phi}^2 + V(\phi) \\ p &= \frac{1}{2}\dot{\phi}^2 - V(\phi), \end{aligned}$$

and $w = \frac{p}{\rho} = \frac{\frac{1}{2}\dot{\phi}^2 + V(\phi)}{\frac{1}{2}\dot{\phi}^2 - V(\phi)} \to -1$. The damping happens because the universe is expanding and tends to damp out relative motion. (It's a scalar, though, and there is no obvious relative motion.)

We call the approximation that the field is overdamped the "slow roll" approximation. This means that we can ignore the $\ddot{\phi}$ term and the kinetic energy density term $\frac{1}{2}\dot{\phi}^2$. This gives

$$\begin{aligned} H^2 &= \frac{\rho}{3\bar{m}^2} = \frac{V(\phi)}{3\bar{m}^2} \\ \dot{\phi} &= -\frac{V'(\phi)}{3H}. \end{aligned}$$

To simplify it even more, we define the "number of e-foldings" $N \equiv \ln\left(\frac{a_{\text{end}}}{a}\right)$. Then $dN = -d\ln a = -H dt$. Then

$$\frac{d\phi}{dN} = -\frac{d\phi}{H dt} = -\frac{\dot{\phi}}{H} = \frac{V'(\phi)}{3H^2} = \bar{m}^2\frac{V'(\phi)}{V(\phi)}$$

or, in general multifield inflation,

$$\frac{d\phi}{dN} = \bar{m}^2 \nabla \ln V.$$

To determine when the slow-roll approximation fails, we define two parameters $\epsilon \equiv \frac{\bar{m}^2}{2}\left(\frac{V'}{V}\right)^2$ and $\eta \equiv \bar{m}^2\frac{V''}{V}$ which need to have magnitudes small compared to unity. During slow roll, we have

$$N = \frac{1}{\bar{m}^2}\int \frac{V(\phi)}{V'(\phi)}d\phi.$$

3.3 Interpretation of the potential $V(\phi)$

In the (local) homogenous and isotropic approximation, we have a function $\phi(\vec{r}, t)$ which does not depend on $\vec{r}$. A constant V never stops inflating. A linear V hits zero and bad things happen.

The next simplest one is $V(\phi) = m^2\phi^2$, which works pretty well. The first slow roll parameter is

$$\epsilon \equiv \frac{\bar{m}^2}{2}\left(\frac{V'}{V}\right)^2 = \frac{\bar{m}^2}{2}\left(\frac{2}{\phi}\right)^2 = \frac{2\bar{m}^2}{\phi^2}$$

and inflation ends when $\phi \sim \phi_{\text{end}} = \sqrt{2}\bar{m}$ when the kinetic energy of the field becomes significant. The second is

$$\eta \equiv \bar{m}^2\left|\frac{V''}{V}\right| = \frac{2\bar{m}^2}{\phi^2}$$

which gives the same result. The number of e-foldings is

$$N = -\frac{1}{\bar{m}^2}\int_{\phi}^{\phi_{\text{end}}}\frac{V}{V'}d\phi = \frac{\phi^2 - \phi_{\text{end}}^2}{2\bar{m}^2}$$

and

$$\phi(N) = \sqrt{\phi_{\text{end}} + 4\bar{m}^2N} = \sqrt{2}\bar{m}\sqrt{1+2N}.$$

3.4 Evidence for inflation

There are multiple approaches to measuring h and the Ω parameters and they are all consistent with the predictions from inflation. The physically meaningful quantity is $\frac{\rho}{\rho_{\text{end}}} = 1 + 2N$, which has no $\bar{m}$. To match the CMB fluctuations, we'd have to be around 10^{-3} of the Planck scale. These are all calculated for ϕ when the horizon was the same size it is now. This will become more clear with the problem set and later lectures. (Problem set hint: $n_s \approx 1 - \frac{2}{N} \approx 0.96$.) Inflation would also produce gravitational waves which we can potentially see. These have wavelengths comparable to the horizon and they should imprint a signal on the polarization of the microwave background.

The parameter $w = \frac{p}{\rho}$ is based on the idea that we don't know what dark energy is. If we assume it follows a power law, then w is constant and we can fit it. The result is consistent with -1.

3.5 Crazy effects of inflation

3.5.1 The landscape view

String theory apparently predicts a whole landscape of possible local theories depending on what minimum of whatever fields we're sitting in. As an example, if we have some wavy potential for a single scalar field, we might have slow-roll inflation forever, inflation followed by recollapse, etc. It turns out that various inflation models make testable predictions, and some are wrong (e.g. ϕ^4).

Parallel universes are not a theory but predictions of certain theories. If you believe in inflation, then you should believe in all of its predictions.

It turns out that slow-roll inflation is forbidden in the most clearly understood class of compactifications of string theory.

In any landscape, inflation takes those bits that can inflate and makes them big, so they actually happen.

3.5.2 Eternal inflation

Inflation generically goes on forever. Most of the universe could get stuck in a local minimum and then inflation could end by quantum tunneling. The part where inflation continues grows faster than it decays away.

If you start with an FRW metric stuck in a local minimum, then any comoving bit will eventually "decay" and stop inflating but the bits in between gain volume even faster.

It turns out that if you just have a local maximum, then the bit on top will inflate for so long that the total (physical) volume increases without bound. When inflation ends, the coordinates tilt so that infinite space comes from what used to be the infinite time coordinate.

Even ϕ^2 inflates eternally. The inflaton field jiggles a little due to quantum mechanics and it turns out that the total volume where you shift up can increase faster than slow roll eliminates it. There's a boundary called the quantum diffusion boundary above which you are driven up more than down. Below that are the slow roll and reheating regimes.

3.5.3 Measure problem and the A-word

Suppose we want to compute the probability that a random proton has some parameters. If we knew the prior distribution (i.e., what a randomly-chosen proton would be) and the probability that a randomly selected proton end up being in one of us conditioned on its parameters, we could compute the a posteriori distribution of what we should expect to see.

We get $f(\bar{p}) \propto f_{\text{prior}}(\bar{p})\, f_{\text{people here}}(\bar{p})$. Both factors are hard to compute, but *we cannot ignore the second one.*

Suppose we want to compute the probability with which some numbers have the values they are observed to have. There are few reasons they could have the values they have. For example, we don't expect the Lagrangian of nature to let us compute the mass of earth or the semimajor axis. We'd like to think that the Lagrangian would let us compute the mass of the electron and the Bohr radius of the hydrogen atom.

There are a few parameters that predict most things: α, the fine structure constant; β, the ratio of electron and proton masses, and m_p, the proton mass in Planck units. These are somewhat coarsely-tuned—bad things happen if they're much different. But there is other fine-tuning: we need the neutron to be heavier than the proton, but not much heavier—otherwise the neutron wouldn't be stable in a nucleus, either. This is due to detailed differences between the quark masses, which requires a rather precise Higgs VEV (in those models, anyway). Factors of order unity blow everything up. Even smaller changes dramatically alter the carbon and oxygen yield in stars.

3.6 Reheating

During inflation, the energy density of the universe is dominated by the potential energy of the field. Once slow roll ends, it turns out that $\rho \propto a^{-3}$, just like ordinary matter. Somehow, the energy left over from inflation must have turned into other stuff. The details of how this works is unknown.

4 Big bang nucleosynthesis

4.1 Basics

Looking backwards in time, the universe was a lot hotter and denser than it is now, and far enough back (even before the last scattering surface), it would have been as hot as in the core of the sun, so nuclear fusion should have happened. It would be a short-lived process because the universe was expanding, so there's a nonequilibrium system to solve. George Gamow did a simple approximation of it and got a result rather close to the modern answer. The amount of helium you end up with depends on only one number, the baryon-to-photon ratio. and it does not depend very strongly. You can measure that ratio from the CMB and compute the abundances of deuterium, helium, helium-3, and lithium-7. (Gamow thought it was a failure b/c astronomers had said that stars couldn't make the heavy elements, which turned out to be wrong. There's a resonance that allows three alpha particles to form carbon.)

There is a certain amount of uncertainty in the abundance of stuff, since we need to correct for post-big-bang nucleosynthesis (in stars), but the error bars are supposed to take that into account and the model seems to work very good.

We know the number and mass density of photons:

$$\begin{aligned} n_\gamma &= \frac{2\zeta(3)}{\pi^2} T^3 \\ \rho_\gamma &= \frac{\pi^2}{15} T^4 \end{aligned}$$

We use T from the CMB (now). We find the baryon-to-photon ratio to be $\eta = \frac{n_b}{n_\gamma} = 6 \cdot 10^{-10}$ (observed). [Note to reader: Max says that no process after freezeout would have changed η. I don't see why.] The evolution of the abundances of species is $\dot{n}_i = -\gamma_i n_i + \sum_{jk} \gamma^i_{jk} n_j n_k$, which is a non-linear stiff system, but we still know how to solve it. We can approximate this by finding the very fast processes and assuming that they are always in equilibrium. Some timescales in the problem are $\Gamma = n \langle \sigma v \rangle$, the rate at which a particle interacts with things, and H, the rate at which densities drop due to expansion. Eventually, the reaction rates drop below the Hubble time and reactions slow down. We can approximate this by equilibrium followed by freezeout.

We start out with the same initial number density of everything, when the temperature is higher than, say, 1 GeV. We'd expect most of the matter and antimatter to anihiliate at low energy, and, in fact, we'd expect to have even less left than we have now. The parameter η has no business being a free parameter—we'd like it to come out of the theory.

4.2 Thermal equilibrium

Some units, first: 1eV = 11604K and 1MeV $\sim 10^{10}$K. 1MeV is around the scale of the mass difference between protons and neutrons, the scale of electron-positron pair production, and the scale of nuclear reactions. Therefore, in the early universe, very little of what's around depends on historical accidents. The number density at equilibrium is

$$n(\bar{p})\, d^3p = \frac{d^3p}{(2\pi^3)} \frac{g}{e^{E/kT} \pm 1},$$

with $E = \sqrt{m^2 + p^2}$, where g depends on the number of degrees of freedom.

species	g	g_*
photon	2	2
neutrinos (left-handed)	6	$\frac{21}{4}$
e^+ and e^+	4	$\frac{7}{2}$

The density of particles is $\rho = 4\pi \int n(p)\, Ep^2 dp$. With $m \ll T$, we get $\frac{g}{2} a_B T^4 = \frac{\pi^2}{30} g T^4$ for bosons and $\frac{7}{8}\frac{\pi^2}{30} g T^4$ for fermions. We define $g_* = \frac{7}{8} g$ for fermions, since all we usually care about is the density. The pressure element is $\frac{p^2}{3E} d^3p$, so the total pressure is $\frac{1}{3}$ the total density. The entropy is $S = \frac{P+\rho}{T} \propto T^3 \propto n$.

Massless particles cool, with $T \propto a^{-1}$ (mean energy just redshifts). Conservation of entropy per unit comoving volume means that $sa^3 \propto g_* T^3 a^3$ is constant.

4.3 Important temperatures

Somewhere around 10MeV, the heavy leptons have decayed and the protons are formed but irrelevent. Neutrinos freeze out around 1MeV (due to the reactions that form them becoming slow). Slightly after neutrino freeze-out, big bang nucleosynthesis occurs. Protons don't matter for cosmology until matter-radiation equality, which happens around 1eV / 10000K. Recombination happens around 3000K.

- At first, $T_\gamma = T_{e^+} = T_\nu$ because of $e^+ + e^- \leftrightarrow \nu + \bar{\nu}$.

- Then neutrinos decouple, but the temperatures track each other because they cool at the same rate.

- Then the remaning electron-positron pairs annihilate. By conservation of entropy,

$$\frac{(aT)^{\text{after}}}{(aT)^{\text{before}}} = \left(\frac{g_*^{\text{before}}}{g_*^{\text{after}}}\right)^{1/3} = \left(\frac{2+\frac{7}{8}4}{2}\right)^{1/3} = \left(\frac{11}{4}\right)^{1/3}.$$

  We can solve this for today's conditions by

$$\frac{T_\nu}{T_\gamma} = \frac{aT_\nu}{aT_\gamma} = \frac{aT_\gamma^{\text{before}}}{aT_\gamma^{\text{after}}} = \left(\frac{4}{11}\right)^{1/3}.$$

  To solve for conditions before matter-radiation equality, we can use

$$\begin{aligned}\rho_{\text{radiation}} &= \rho_\gamma + \rho_\nu \\ \rho_\nu &= \frac{\pi^2}{30}\cdot\frac{21}{4}T_\nu^2 = \rho_\gamma\cdot\frac{21}{8}\left(\frac{4}{11}\right)^{4/3},\end{aligned}$$

  which is a large correction to the number of photons. We can try to measure the prefactor to try to count neutrinos. This depends on freezeout temperatures for hypothetical extra neutrinos, but it still allows some models to be ruled out. (The theoretical number of neutrinos is $N_\nu \approx 3.04$ due to a little extra heat leaking out into neutrinos.)

4.4 BBN Reactions

$$\begin{aligned}\mathrm{n}+\mathrm{e}^+ &\leftrightarrow \mathrm{p}^+ + \bar{\nu} \\ \mathrm{n}+\nu &\leftrightarrow \mathrm{p}^+ + \mathrm{e}^- \\ \mathrm{n} &\leftrightarrow \mathrm{p}^+ + \mathrm{e}^- + \nu\end{aligned}$$

The first two reactions have rate $\frac{\Lambda}{H} \approx \left(\frac{T}{0.8MeV}\right)^3$. The T^3 comes from $\Lambda = \underbrace{n}_{T^3}\langle\underbrace{\sigma}_{T^2}\underbrace{v}_{1}\rangle$. To solve these things, we need to know the cross-sections, temperatures, and densities. We know the first two, and we know the number density of photons, both before and after BBN. We know the pre-1MeV ratios of the baryon species, but we don't know the total number, because the baryons came from physics we don't know. The only number we need to find all of these is $\eta = \frac{n_b}{n_\gamma}$.

4.5 WIMPs

A WIMP is a weakly interactive massive particle. We can't have strongly interacting massive particles (easy to detect with mass spectroscopy, for example) or electromagnetically interacting massive particles (easy to see). We know that dark matter interacts gravitationally, and if the particles are not weakly interacting (GIMPs), then we're screwed.

One redeeming feature of WIMPs is that a simple estimate gives a decent estimate of the amount of dark matter. In Planck units ($c = \hbar = G = 1$), where $\xi_{\mathrm{WIMP}} = \frac{m_{\mathrm{WIMP}}^3}{g^4 T}$ is the density of WIMPs in Planck masses per photon,

$$\begin{aligned} \langle \sigma v \rangle &\sim \frac{\alpha_w^2}{M_w} \\ n_{\mathrm{WIMPS}} &= \frac{\rho_{\mathrm{WIMP}}}{M_{\mathrm{WIMP}}} = \frac{n_\gamma \xi_w}{M_{\mathrm{WIMP}}} \\ \text{Annihilation rate } \Gamma &\approx \langle \sigma v \rangle \, n_{\mathrm{WIMP}}, \end{aligned}$$

giving $\xi_{\mathrm{WIMP}} \sim 10^5 \frac{v^2}{g^4} \sim 10^{-28}$. This is close to the measured value.

5 The universe to 1st order

Einstein's equations are nonlinear PDEs and they're messy to solve directly. But the perturbations we have are small, though (initially on the order of 10^{-5}), so we can linearize and Fourier transform.

5.1 Newtonian analysis

We have three equations.

Conservation of mass $\dot{\rho} + \nabla \cdot (\rho \vec{v}) = 0$

Euler EQ $\dot{\vec{v}} + \nabla \vec{v} = -\nabla \left(\phi + \frac{p}{\rho} \right)$

Poisson EQ $\nabla^2 \phi = 4\pi G \rho$

These hold approxmately in GR, too, due to Birkhoff's Theorem, which says that, in sperically symmetric space, things at higher radius don't matter. So we can ignore GR.

We now expand around a zeroth-order solution.

$$\begin{aligned} \rho_0 (t, \vec{r}) &= \bar{\rho}(t) \propto a^{-3} \text{ if MD} \\ \vec{v}_0 (t, \vec{r}) &= \frac{\dot{a}}{a} \vec{r} \\ \phi (t, \vec{r}) &= \frac{2\pi G \bar{\rho}}{3} r^3 \end{aligned}$$

The traditional expansion is $\rho (t, \vec{r}) = \bar{\rho} [1 + \delta (t, \vec{r})]$, $\vec{v} (t, \vec{r}) = \vec{v}_0 (t, \vec{r}) + \vec{v}_1 (t, \vec{r})$, and $\phi (t, \vec{r}) = \phi_0 (t, \vec{r}) + \phi_1 (t, \vec{r})$. When we plug these in, the zeroth-order terms cancel by construction, and we ignore the second-order terms (they are smaller by five orders of magnitude), so we keep the first-order terms and end up with linear equations.

To find $\bar{\rho}$, let $w \equiv \frac{p}{\rho}$ be the equation of state. If we expand a box a little bit, then thermodynamics says $dE = -pdV = -w\rho dV$. Relativity says

$$\begin{aligned} d\rho &= d\left(\frac{E}{V}\right) \\ &= \frac{dE}{V} - \frac{E}{V^2}dV \\ &= -\frac{w\rho dV}{V} - \frac{\rho}{V}dV \\ &= -\rho\left(w+1\right)dV \\ d\ln\rho &= -\left(1+w\right)d\ln V \end{aligned}$$

If w is constant, then $\rho \propto V^{-(1+w)} \propto a^{-3(1+w)}$.

The result is $\ddot{\hat{\delta}}(k) + 2H\dot{\hat{\delta}}(k) + \left(\frac{v_s^2 k^2}{a^2} - 4\pi G\bar{\rho}\right)\hat{\delta}(k) = 0$. v_s is the speed of sound. In the case of zero pressure, there is no k dependence.

In the (inconsistent) case of a static universe, there are oscillatory solutions, a runaway solution, and a decay solution. The sign depends on whether pressure or gravity wins. The cutoff is $\frac{k}{a} = \frac{\sqrt{4\pi G\bar{\rho}}}{v_s} \to \lambda_J = \frac{2\pi a}{k} = \sqrt{\frac{\pi}{G\bar{\rho}}}v_s$, the Jeans wavelength. One way to guess this is to compare the $\tau_G \sim \frac{1}{\sqrt{G\rho}}$, the timescale of gravitational collapse with $\tau_P \sim \frac{\lambda}{v_s}$, the timescale of pressure waves. If the gravitional timescale is smaller, then we have instabilities.

With no curvature and no pressure, we plug in the Friedman equation and get $\ddot{\delta} + 2H\dot{\delta} - \frac{3}{2}H^2\delta = 0$. This is bizarre—the overdensities do not depend on their values far away. Assuming a MD universe, $H \propto \frac{2}{3}t^{-1}$, so $\ddot{\delta} + \frac{4}{3}t^{-1}\dot{\delta} - \frac{2}{3}t^{-2}\delta = 0$. Making the ansatz $\delta = At^n$ and solving the resulting quadratic equation gives $n = -\frac{1}{6} \pm \frac{5}{6} \in \left\{-1, \frac{2}{3}\right\}$. So $\delta \propto a$ or $\delta \propto a^{-3/2}$, and there is a growing mode and a decaying mode. The decaying mode decays away, so the moral of the story is

$$\delta \propto a.$$

This is modified by a few things. First, modes larger than the horizons don't amplify because no information propagates. Second, in a radiation-dominated universe, amplificaiton behaves differently. Before matter-radiation equality, fluctuations don't grow much. After matter-radiation equality, growth speeds up. Finally, when $\delta \sim 1$, nonlinear effects prevail.

5.2 Measuring the power spectrum

We imagine the variance of something being related to some distribution. But the universe we isn't a random variable—it's a single instance of one. It turns out the entire probability distribution for the overdensity is characterized by a single function of a single variable.

A Gaussian random vector is a vector with joint density

$$f(\vec{x}) = \frac{1}{(2\pi)^{1/2}|C|^{1/2}} \exp\left(-\frac{1}{2}\left(\vec{x}-\vec{\mu}\right)^T C\left(\vec{x}-\vec{\mu}\right)\right).$$

Because the harmonic oscillator ground state is a Gaussian, we expect a Gaussian distribution for the overdensity.

We can characterize the distribution for an arbitrary random field by giving the entire family of n-point functions

$$f_n\left(\delta\left(\vec{r_1}\right),\ldots,\delta\left(\vec{r_n}\right)\right),$$

which depend only on the mean and covariance. For the overdensity, we've defined the mean to be zero. The second moment $C_{ij} = \langle\delta_i,\delta_j\rangle = \langle\delta\left(\vec{r_i}\right),\delta\left(\vec{r_j}\right)\rangle = \xi\left(|\vec{r_i}-\vec{r_k}|\right)$ because the universe is homogenous and isotropic.

We define the Fourier-space overdensity

$$\begin{aligned}
\delta\left(\vec{r}\right) &= \frac{1}{(2\pi)^3}\int e^{i\vec{k}\cdot\vec{r}}\hat{\delta}\left(\vec{k}\right)d^3\vec{k} \\
\hat{\delta}\left(\vec{k}\right) &= \int e^{i\vec{k}\cdot\vec{r}}\delta\left(\vec{r}\right)d^3\vec{r} \\
\left\langle\hat{\delta}\left(\vec{k}\right)^*\hat{\delta}\left(\vec{k}'\right)\right\rangle &= (2\pi)^3\,\delta\left(\vec{k}-\vec{k}'\right)\int e^{-i\vec{a}\cdot\vec{k}}\xi\left(\vec{a}\right)d^3a \\
& (2\pi)^3\,\delta\left(\vec{k}-\vec{k}'\right)\hat{\xi}\left(\vec{k}\right),
\end{aligned}$$

et voila! $P\left(\vec{k}\right) = \hat{\xi}\left(\vec{k}\right)$ is called the power spectrum. Integrating out the angular part gives a decomposition by spherical Bessel functions $P(k) \doteq \int_0^\infty \xi(r)\,j_0(kr)\,dr$. There's a dimensionless version that is the same in all books $\Delta^2(k) = \frac{4\pi}{(2\pi)^3}k^2P(k)$. If X is the average overdensity with some window function $\hat{w}(k)$, then $\left\langle X^2\right\rangle = \int\Delta^2(k)\,|\hat{w}(k)|\,d\ln k$. So Δ^2 gives the variance if smoothed out on a scale of k. (Yes, this is vague. Maybe I'll fix it.)

6 The transfer function $T(z,k)$

6.1 Multi-component perturbation growth

The universe contains a number of types of things.

Neutrinos They have a little mass, so they're significant until VD (around now).

CDM Cold dark matter didn't matter much until matter-radiation equality but matters a lot until VD.

Dark energy Irrelevant until around now.

Normal matter After recombination, they were like CDM. But before recombination, there was a photon-baryon fluid, which could have done weird things between matter-radiation equality and recombination. This fluid's pressure was dominated by photons (so the speed of sound was around $\frac{c}{\sqrt{3}}$) but its density had important contributions from both.

Photons After recombination, they stream freely, but their density doesn't matter for very long.

Because there are multiple components, they each feel only their own pressure, but they feel gravity from each other. The new equation is

$$\ddot{\delta}_i + 2\frac{\dot{a}}{a}\dot{\delta}_i + (c_i)^2 k^2 \delta_i - 4\pi G \sum_j \bar{\rho}_j \delta_j = 0.$$

In the matter-dominated era, $\delta \propto a$. In the VD late universe, growth stops. In the the RD era, most of the energy was in photons, which wouldn't notice dark matter clumping. In the photon-baryon fluid, $\lambda_{\text{Jeans}} \sim \lambda_{\text{horizon}}$, so no growth.

We can solve the rolloffs analytically. For RD and MD, $G = 1 + \frac{3}{2}\frac{a}{a_{\text{eq}}}$. In the late universe, if we define $x \equiv \frac{\rho_R}{\rho_M} = \frac{\Omega_\Lambda}{\Omega_M}(1+z)^{-3} \propto a^3$, then $G_\Lambda(x) = \frac{5}{6}\sqrt{1+\frac{1}{x}}\int_0^x \frac{dy}{y^{1/6}(1+y)^{3/2}}$, which is reasonably well approximated by $x^{1/3}\left[1+\left(\frac{x}{G_\infty^3}\right)^\alpha\right]^{-\frac{1}{3\alpha}}$, where $\alpha \approx 0.795$ and $G_\infty = \frac{5\Gamma\left(\frac{2}{3}\right)\Gamma\left(\frac{5}{6}\right)}{3\sqrt{\pi}} \approx 1.437$. So we get an extra $\sim 44\%$ of growth during VD and an extra $\frac{5}{2}$ during RD. By matching and combining the approximations, we get $G(x) = 1 + \frac{3}{2}x_{\text{eq}}^{-\frac{1}{3}} G_\Lambda(x)$.

In the photon-baryon fluid, there were driven oscillations, and then, around recombination, something called Silk damping happened that destroyed many of these fluctuations, a result of the fact that recombination wasn't instantaneous—there was a period where the mean free path of the photons is long but not infinite.

The horizon scale $a_H \propto ct \propto (1+z)^{-3/2}$ when matter dominated. The physical wavelength $a_p \sim \frac{\lambda}{1+z}$. A mode enters the horizon when $a_H = a_p \iff (1+z)^{-3/2} \propto \frac{\lambda}{1+z}$. Then $1+z_{\text{entry}} \propto \lambda^{-2} \propto k^2$. $\delta \propto \Delta = \Delta_{\text{enter}} \begin{cases} 1+z_{\text{enter}} & \text{if } z_{\text{enter}} < z_{\text{eq}} \\ 1+z_{\text{eq}} & \text{otherwise} \end{cases}$. So Δ increases like k^2 up to some value k_{eq}, when it flattens out. ($\Delta^2 \propto k^3 p$).

Clarification: Δ, P, and G are not random. δ is random. G is how much the mode gets multiplied by, once it's in the horizon. The power spectrum at any time can be factored $P(k,z) = P_*(k)\,G(k,z)^2$, where P_* is the very early universe power spectrum, and G is the amplification part.

6.2 Neutrinos

The effect of neutrinos is to suppress the growth of fluctuations on a small scale. There is a characteristic scale; well above the scale there is no suppression and well below it there is suppression by a constant factor. The scale depends on the neutrino mass.

If only a fraction Ω_* of the matter can cluster, then growth decreases to $\delta \propto a^p$, where $p = \left[\sqrt{1+24\Omega_*} - 1\right]/4 \approx \Omega_*^{3/5} \approx (1-f_\nu)^{3/5}$.

Neutrinos don't cluster on small scales, because they are above the escape velocity of a clump. The net growth today is $\sim 4700^{p} \approx 4700 e^{-4f_\nu}$, so the power suppression $P(k)/P(k)_{\text{no neutrinos}} \sim e^{-8f_\nu}$. $f_\nu = \sum m_\nu{}^i / 94.4eV\omega_{dm}$.

Neutrinos can be distinguished from curvature or dark matter in that they actually change the shape of the curve. With enough resolution, the masses of all three species of neutrinos could be read off.

7 The cosmic microwave background

7.1 Observations

Because we're looking at the inside surface of a sphere, we decompose the temperature in spherical harmonics:

$$\frac{\delta T}{T}(\hat{r}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} a_{lm} Y_{lm}(\hat{r}).$$

We expect isotropy, so $\langle a_{lm}^* a_{l'm'} \rangle = \delta_{ll'} \delta_{mm'} c_l$, where c_l is the variance. The normalization gives

$$\frac{\delta T_l}{T} \equiv \sqrt{\frac{l(l+1)}{2\pi} c_l}.$$

To estimate the variance, one way is to look at the values for each m, which are uncorrelated, and use the usual variance estimator. With a perfect detector, $\frac{\Delta c_l}{c_l} = \sqrt{\frac{2}{2l+1}}$.

John Mather showed that the CMB was, in fact, a blackbody. COBE took 7 degree resolution pictures at three frequencies.

After COBE, people took higher-resolution pictures from high-altitude balloons.

WMAP measures the difference between one position and another position 60 degrees away. Then they invert the matrix.

It turns out that the polarization of the CMB might be very interesting. This is difficult to measure, because the polarization signal is around 1% of the fluctuations, and the synchotron emission from the galaxy may be strongly polarized.

7.2 The CMB fluctuation formula

$$\frac{\Delta T}{T}(\hat{r}) = \phi(\vec{r}) - \hat{r} \cdot \vec{v}(\vec{r}) + \frac{1}{3}\delta(r)$$

In this equation, $|\vec{r}|$ is the distance to last scattering.

Why are there fluctuations? There are three answers. The photons are redshifted on the way out of clumps. There is a Doppler effect due to velocities. The density of photons is $\rho \propto T^4$, so overdense clumps are hotter.

Why are there wiggles? The first bump (around 0.5°) is due to the speed of sound—there are sound waves flying around. The scale is just the horizon size at recombination.

What we actually see is a combination of:

- Primary (Effects from $z \gtrsim 10^3$)
  - Gravity, doppler, and density (above)
  - Damping
  - Topological defects, maybe
- Secondary
  - Gravity
    - Early integrated Sachs-Wolfe effect (ISW)
    - Late ISW
    - Rees-Sciama effect
    - Lensing
  - Reionization
    - Local (SZ)
      - Thermal
      - Kinematic
    - Global
      - Suppression
      - New doppler
      - Ostriker-Vishniac
- Tertiary
  - Foregrounds
  - Headaches

If the universe is MD and linear, then $\dot{\phi} = 0$. This means that secondary redshifting due to gravity doesn't happen unless MD or linearity is violated. Sachs and Wolfe derived a formula $\frac{\delta T}{T} = \int \dot{\phi}\,(\vec{r}(t))\,dt$. The early ISW effect is due to radiation, and the late ISW effect is due to dark energy or curvature. The late ISW only started to matter recently, so it would affect the largest scales. The early ISW boosts the first peak, because that was around the horizon size at last scattering.

Rees-Sciama is $\dot{\phi}$ due to structure formation. This shows up at the angular scale of nonlinear structures. In practice, this is a couple of orders of magnitude down from the primary signal, and no one has seen it.

CMB lensing tends to smooth out the power spectrum.

There are also projection effects due to the fact that the power spectrum is 3D but we're looking at a spherical surface. The reinization surface wasn't perfectly sharp, so the last scattering surface has some approximately Gaussian width, which smooths the 2D map and smooths it in the radial direction. This is a low-pass filter and kills off the high-frequency components. The scale of this is around the thickness of the last scattering surface.

7.3 Spherical harmonics

The spherical harmonics are

$$Y_{lm}\left(\hat{r}\right) = \sqrt{\frac{2l+1}{2\pi}\frac{(l-m)!}{(l+m)!}}P_l^m\left(\cos\theta\right)e^{im\phi}.$$

They obey a bunch of formulae:

$$\int Y_{lm}Y^*_{l'm'}d\Omega = \delta_{ll'}\delta_{mm'},$$

$$\sum_{lm} Y^*_{lm}\left(\hat{r}\right)Y_{lm}\left(\hat{r}'\right) = \delta\left(\hat{r},\hat{r}'\right)$$

$$\sum_{m=-l}^{l} Y^*_{lm}\left(\hat{r}\right)Y_{lm}\left(\hat{r}'\right) = \frac{2l+1}{4\pi}P_l\left(\hat{r}\cdot\hat{r}'\right)$$

$$\int P_l\left(\hat{k}\cdot\hat{r}\right)P_{l'}\left(\hat{k}\cdot\hat{r}'\right)d\Omega_k = \left(\frac{4\pi}{2l+1}\right)\delta_{ll'}P_l\left(\hat{r}\cdot\hat{r}'\right)$$

$$e^{i\mathbf{k}\cdot\mathbf{r}} = \sum_{l=0}^{\infty}\left(2l+1\right)i^l P_l\left(\hat{k}\cdot\hat{r}\right)j_l\left(kr\right)$$

We can use these to study correlations:

$$\begin{aligned}
\frac{\delta T}{T}\left(\hat{r}\right) &= \sum_{lm} Y_{lm}\left(\hat{r}\right)\\
\left\langle a^*_{lm}a_{l'm'}\right\rangle &= \delta_{ll'}\delta_{mm'}c_l\\
c\left(\cos\theta\right) &= c\left(\hat{r}\cdot\hat{r}'\right)\\
&= \left\langle \frac{\delta T}{T}\left(\hat{r}\right)^*\frac{\delta T}{T}\left(\hat{r}'\right)\right\rangle\\
&= \left[\sum_{lm} a^*_{lm}Y_{lm}\left(\hat{r}\right)^*\right]\left[\sum_{lm} a_{lm}Y_{lm}\left(\hat{r}'\right)\right]\\
&= \sum_{lml'm'}\left\langle a^*_{lm}a_{l'm'}\right\rangle Y_{lm}\left(\hat{r}\right)Y_{l'm'}\left(\hat{r}'\right)\\
&= \sum_{lm} c_l Y^*_{lm}\left(\hat{r}\right)Y_{lm}\left(\hat{r}'\right)\\
&= \sum_{l=0}^{\infty}\left(\frac{2l+1}{4\pi}\right)P_l\left(\hat{r}\cdot\hat{r}'\right)
\end{aligned}$$

7.4 Projection effects

The effect of a source being in a gravitational well is the redshift, giving $\frac{\delta T}{T} = \phi$, except that things in a gravitational well cool more slowly (due to time dilation), which means it's hotter, which has the opposite sign. $T \propto a^{-1} \propto t^{-2/3}$, so $\frac{\delta T}{T} = -\frac{2}{3}\frac{\delta T}{T} = -\frac{2}{3}\phi$, giving an actual contribution of

$$\frac{\delta T}{T} = \frac{1}{3}\phi.$$

Now we write

$$\begin{aligned}
\frac{\delta T}{T}(\hat{r}) &= \frac{1}{3}\phi(\hat{r}) \\
&= \frac{1}{3}\frac{1}{(2\pi)^3}\int e^{i\mathbf{k}\cdot\mathbf{r}}\hat{\phi}(\mathbf{k})\,d^3\mathbf{k} \\
c(\hat{r}\cdot\hat{r}') &= \left\langle \frac{1}{9}\frac{1}{(2\pi)^6}\left[e^{-i\mathbf{k}\cdot\mathbf{r}}\hat{\phi}^*(\mathbf{k})\,d^3\mathbf{k}\right]\left[e^{i\mathbf{k}'\cdot\mathbf{r}'}\hat{\phi}^*(\mathbf{k}')\,d^3\mathbf{k}'\right]\right\rangle \\
&= \frac{1}{9}\frac{1}{(2\pi)^6}\iint e^{-i\mathbf{k}\cdot\mathbf{r}}e^{i\mathbf{k}'\cdot\mathbf{r}'}\underbrace{\left\langle\hat{\phi}^*(\mathbf{k})\,\hat{\phi}(\mathbf{k}')\right\rangle}_{=(2\pi)^3\delta(\mathbf{k}-\mathbf{k}')P_\phi(\mathbf{k})}d^3\mathbf{k}d^3\mathbf{k}' \\
&= \frac{1}{9}\frac{1}{(2\pi)^3}\iint\left(e^{i\mathbf{k}\cdot\mathbf{r}}\right)^* e^{i\mathbf{k}\cdot\mathbf{r}'}P_\phi(\mathbf{k})\,d^3\mathbf{k} \\
&= \frac{1}{9}\frac{1}{(2\pi)^3}\int\left[\sum_l(2l+1)(-i)^l P_l\left(\hat{k}\cdot\hat{r}\right)j_l(kr)\right]\left[\sum_{l'}(2l'+1)\,i^l P_l\left(\hat{k}\cdot\hat{r}'\right)j_l(kr)\right]k^2dkd\Omega \\
&= \frac{1}{9}\frac{1}{(2\pi)^3}\int\sum_l(2l+1)^2 j_l(kr)\,j_l(kr')\,\frac{4\pi}{2l+1}P_l(\hat{r}\cdot\hat{r}')\,P_\phi(\mathbf{k})\,d^2dk \\
&= \frac{1}{9}\frac{(4\pi)^2}{(2\pi)^3}\int\sum_l\frac{(2l+1)}{4\pi}j_l(kl)^2\,P_\phi(\mathbf{k})\,P_l(\hat{r}\cdot\hat{r}')\,k^2dk \\
c_l &= \frac{1}{9}\frac{\left(4\pi^2\right)}{(2\pi)^3}\int j_k(kl)^2\,P_\phi(\mathbf{k})\,k^2dk
\end{aligned}$$

This gives the expected spherical harmonic power as a function of the 3D power spectrum.

7.5 Nonlinear effects

When the linear theory predicts an overdensity of ˜1.69, the nonlinear theory predicts infinite density, which means that the clump "virializes" and forms a bound object with density around 200 times the background.

The Press-Schechter approximation says that a point is in a virialized halo of mass exceeding M if an appropriately smoothed δ at that point exceeded 1.69 at the appropriate time, which gives something like erfc $\left[\frac{\delta_c}{\sigma(M,z)}\right]$, which grows very quickly as the number of sigmas changes. The trouble with interpreting this data is figuring out the mass. We can now look at the clusters in visible light, X-rays, lensing, etc.

7.6 Seeing other effects

7.6.1 ISW

We now see the late ISW effect at respectable significance. We can't see the Rees-Sciama effect yet.

7.6.2 Global reionization

Plasma due to reionization scatters light. We parameterize it by the reionization optical depth $\tau = \int \sigma_T n_e dx$. This washes out large-$l$ power. (For $l \gg 10$, $c_l \sim c_l e^{-2\tau}$.) For smaller l, there is little effect. There is some increase at medium scales due to fluctiations in the plasma. There's some annoyance because changing the seed fluctuation size and τ only affects small l in a manner that's somewhat degenerate with the ISW.

The OV (Ostriker-Vishniac) effect is a second-order effect during reionization that may be apparent at very small scales ($l \sim 10^4$). It's like a kinetic SZ effect caused by general diffuse hydrogen.

7.6.3 Local reionization

The thermal SZ effect occurs because galaxy clusters contain hot gasses, at around 5-15keV. This changes the spectrum by heating some cold photons. It's detectable by the frequency spectrum of the CMB. It's a suppression below $V_* = 217$GHz and an amplification above. (V_* is called the SZ null frequency and is a small multiple of the CMB temperature. Relativistic gasses can perturb it 1-2MHz higher.) The suppression in antenna temperature asymptotes to a constant at low frequency (i.e. constant decrease in T). At higher frequencies, $\Delta T \propto V - V_*$.

The strength of the SZ effect is independent of cluster redshift, so we can find all clusters by looking for holes in the CMB.

Galaxy clusters are also moving, which causes them to Doppler shift the CMB, about 10 times less than the thermal SZ effect. We can use this to measure the radial velocities of the clusters by changes at the SZ null frequency. (This is the same physics as the Ostriker-Vishniac effect.)

7.7 CMB polarization

The last thing that the CMB light did before reaching us was to Thompson scatter off an electron. Thompson scattering is highly polarizing, but this effect is neutralized because light hits scattering electrons from all directions. If there is quadrupole anisotropy in the CMB, though, there will be some residual polarization.

7.8 Measuring the CMB

At low frequencies (˜23GHz), we see synchotron radiation from the galaxy. At high frequencies, we see blackbody emission from dust in the galaxy.

The WMAP team's approach to removing this foreground is to take a linear combination of the five maps (at different frequencies) that minimizes total power. If the weights are allowed to depend on the multipole moment, then you can do better, because different types of foregrounds are dominant at different angular scales.

Polarization is harder because the signal is weaker and the foregrounds are stronger (e.g. synchotron radiation could be 70% polarized).

The polarization field can be decomposed into T, E and B modes (related to the angle between the polarization and the gradient). This gives three scalars. The power spectra and cross-power spectra can be measures (it turns out that the cross-power that can be nonzero is TE). B is not expected to be amplified during expansion, but inflation can produce B modes. The Sachs-Wolfe effect in particular does not polarize at all.

8 Lensing

Gravity bends light. If we see a point source lined up with a point mass, it'll show up as a ring, and, if it's a little off, we'll see an arc and some points instead. Farther out, things look more elliptical than they are. If we look far away and see correlated ellipticities, then this might indicate dark matter (it turns out that nearby galaxies might have correlations due to tides).

Lensing has the property that the surface brightness (flux per unit solid angle) never changes.

In this section, unprimed coordinates are pre-lensing, and primed coordinates are what we see in the presence of lensing.

$$\begin{aligned} F' &= \int \phi'(\vec{\theta'}) d^2\vec{\theta'} \\ &= \int \phi'(\vec{\theta'}) \left| \frac{d\vec{\theta'}}{d\vec{\theta}} \right| d^2\theta \end{aligned}$$

The Jacobian is nearly constant over the support of the flux (i.e. the size of the galaxy), so

$$F' = F\,|\mathbf{J}|\,.$$

This is annoying because we don't know F in general. We could do a similar calculation for the position shift, but we have the same problem because we don't know where the source really is.

If we define $f(\theta') = \frac{\phi(\theta')}{F}$ to be the normalized density of where the light comes from, then the second moment is the quadrupole (covariance) $\mathbf{Q} =$

$\left\langle \left(\vec{\theta} - \left\langle \vec{\theta} \right\rangle\right) \left(\vec{\theta} - \left\langle \vec{\theta} \right\rangle\right)^T \right\rangle$. We can linearize the mapping, giving $\vec{\theta}' = \left\langle \vec{\theta}' \right\rangle + J\left(\vec{\theta} - \left\langle \vec{\theta} \right\rangle\right)$. The quadrupole simplifies to $\mathbf{Q}' = \mathbf{J}\mathbf{Q}\mathbf{J}^T$. We assume that the universe is isotropic, even in a particular direction, so $\langle \mathbf{Q} \rangle_{\text{galaxies}} = \sigma^2 \mathbf{I}$. The transformation of the quadrupole is linear, so $\langle \mathbf{Q}' \rangle_{\text{galaxies}} = \sigma^2 \mathbf{J}\mathbf{J}^T$.

One way to relate this to the physics is by the time delay $\tau \propto \frac{1}{2}\left(\vec{\theta}' - \vec{\theta}\right)^2 - \psi(\theta)$ where $\psi(\theta)$ is the gravitational time delay. With multiple images, we can measure this directly. But with Fermat's principle, there is an image where $0 = \nabla\tau = \vec{\theta}' - \vec{\theta} - \nabla\psi \iff \vec{\theta}' = \vec{\theta}\nabla\psi$. The Jacobian is thus $\mathbf{J}_{ij} = \delta_{ij} + \psi_{,ij} = \mathbf{I} + \mathbf{M}$, where $\mathbf{M}$ is the Hessian. This means that the Jacobian matrix is symmetric, and we can parameterize it as $\mathbf{J} = \begin{pmatrix} 1+\kappa+\gamma_1 & \gamma_2 \\ \gamma_2 & 1+\kappa+\gamma_1 \end{pmatrix}$. κ is the *convergence* and the γ's are the *shears*. The magnification is $J = (1+\kappa)^2 - \gamma_1 - \gamma_2$. In weak lensing, these numbers are small, so $J \approx 1 + 2\kappa$.

From now on, ϕ is Newtonian gravitational potential.

In Euclidean space, $\delta\vec{\theta} = \int A(r) \nabla_\perp \phi dr$. $A(r)$ is a slowly varying geometrical factor, and, if the lensing is caused by a galaxy cluster at a narrow range of z, $\delta\vec{\theta} \approx \nabla_\perp \left(A(r) \int \phi dr\right) = \nabla_\perp \psi$.

We can now relate lensing to density (Σ is the surface density of the cluster):

$$\begin{aligned} \nabla_\perp^2 \psi &= A \int \left(\nabla^2 - \partial_r^2\right) \phi dr \\ \int \nabla^2 \phi dr &= \int 4\pi G \rho dr = 4\pi G \Sigma \\ \int \partial_r^2 \phi dr &= \phi'(r_{max}) - \phi'(r_{min}) = 0. \end{aligned}$$

So $1 + \kappa = \frac{1}{2}\operatorname{Tr}\mathbf{J} = \frac{1}{2}\nabla_\perp^2\psi = \frac{1}{2}\frac{\Sigma}{\Sigma^2}$.

After accounting for the expanding universe and other such details, the power spectrum of κ is $P_\kappa(l) = \left(\frac{3}{2}\frac{\omega_m}{c}\right)^2 \int_0^{\eta_{\text{hor}}} \bar{W}(\eta)^2 P\left(\frac{l}{S_k(\eta)}, z(\eta)\right)(1 + z(\eta))^2 d\eta$, where $\bar{W}(\eta) = \int_\eta^{\eta_{\text{hor}}} G(\eta') \frac{S_k(\eta'-\eta)}{S_k(\eta')} d\eta'$ and $G(\eta)$ is the distribution of source redshifts, and the popular fit is $G(\eta) = \frac{B}{z_0 \Gamma\left(\frac{1+\alpha}{B}\right)}\left(\frac{z}{z_0}\right)^\alpha \exp\left[-\left(\frac{z}{z_0}\right)^\beta\right]$. The good news is that it shows the location of all matter as opposed to just baryonic matter. The bad news is that there are nonlinear growth effects, blobs of gas, and other stuff. There is also a problem in that anisotropic point spread functions in telescopes looks just like lensing, so some corrections are needed. These effects are large but can be corrected with care. It turns out that you can even do lensing tomography if you measure each galaxy's redshift individually.

9 The Lyman α forest

If we look at a distant quasar, the absorption due to the (redshifted) Lyman α line from neutral hydrogen in the way gives a map of density along that line. There are some issues due to Doppler broadening (temperature-dependent) and ionization, but it has good sensitivity on small scales. The best data is from SDSS, which has spectra for about 10^5 quasars, and the Keck high-res spectra, which has fewer quasars.

10 21-cm tomography

We would like to measure the density field δ in the observable universe. We've covered a very small portion of the volume so far. The density ratio of triplets to singlets is

$$\frac{n_1}{n_0} = 3 \exp\left[T_*/T_s\right]$$

where $T_* =??$. T_S is the spin temperature. Early on, the CMB does a good job of driving the spin transition, so the spin temperature would be the CMB temperature and there would be nothing to see. But gas cools faster than light (a^{-2}) once the gas is decoupled enough, so the hydrogen ends up cooling faster because collisions cool it faster than heating from the CMB. As a result, we'll see it being colder than the CMB. Once reionization starts, the spin temperature will increase due to UV and the hydrogen will heat up dramatically. This means that around $9 < z < 50$ we'll see it in absorption and at $z < 9$ or so we'll see it in emission. The fluctuation in the brightness temperature is around

$$\delta T_b \approx 29mK \left(\frac{h\Omega_b}{0.03}\right)\left(\frac{\Omega_m}{0.25}\right)^{-1/2}\sqrt{\frac{1+z}{10}}\left(1+x\right)\left(1+\delta\right)\left(\frac{T_S - T_{CMB}}{T_S}\right)$$

where x is the ionization fraction.

There are all kinds of interesting models to rule out at very small scales.