Week 10: Theoretical predictions for the CMB temperature power spectrum

April 5, 2012

1 Introduction

Last week we defined various observable quantities to describe the statistics of CMB temperature fluctuations, chief among them the power spectrum $C_l^{\rm TT}$. This week we will discuss the theoretical prediction for the CMB temperature power spectrum from inflationary density perturbations. The precise calculation is straightforward but extremely complicated. Rather than push through in its entirety, we will simply calculate it in the large-angle limit, and only under some simplifying assumptions. We will then discuss some of the qualitative features of the power spectrum on smaller scales. Next week we will discuss CMB polarization and then inflationary gravitational waves. Interestingly enough, the full calculation of the CMB power spectrum is simpler in that case than for the density perturbations because the complications associated with gauge choices with scalar perturbations are avoided, and we will go through that in detail.

To begin, let's return to our earlier discussion of density perturbations from inflation. There we showed that the freezeout of quantum fluctuations in the inflaton gives rise to a primordial spectrum of inflaton (ϕ) perturbations with power spectrum

$$P_\phi(k) = \left(\frac{H}{2\pi}\right)^2\Bigg|_{k=aH}. \qquad (1)$$

We then waved our hands and said that if the inflaton dominates the energy density of the Universe during inflation, then inflaton fluctuations $\delta\phi$ should induce density perturbations $\delta\rho$. We will now make this connection precise, recalling that the definitions of these inflaton and density perturbations is gauge dependent; it requires specification of the coordinates we choose for the perturbed FRW spacetime.

Recall that a coordinate choice on the perturbed FRW involves a *threading*: i.e., a selection of timelike curves of fixed $\vec{x}$. Since the gradients in any such threading (i.e., the physical separation at fixed t between any two points of fixed $\vec{x}$) are already first order in perturbations, the perturbations to scalar quantities (the density and also ϕ) must be independent of our choice of the threading.

A coordinate choice also involves a *slicing* of the perturbed FRW spacetime; i.e., a choice of surfaces of constant t. Recall that the *comoving slicing* is the slicing in which hypersurfaces of constant t

are perpendicular to worldines of comoving (although not necessarily freely-falling) observers. And the *spatially flat* slicing is that in which the spatial curvature ${}^{(3)}R = 0$.

Now recall that the most general perturbed FRW spacetime can be written,

$$ds^2 = a^2(\tau)\left\{-(1+2\psi)d\tau^2 + 2w_i d\tau dx^i + [(1-2\varphi)\gamma_{ij} + 2h_{ij}]dx^i dx^j\right\}, \tag{2}$$

with h_{ij} traceless: $\gamma^{ij}h_{ij} = 0$. Remember also that $\psi(\vec{x},\tau)$ and $\varphi(\vec{x},\tau)$ are scalars, $w_i(\vec{x},\tau)$ is a vector, and $h_{ij}(\vec{x},\tau)$ is a symmetric trace-free tensor. We also choose h_{ij} to be traceless, as any trace can be absorbed into φ. It is straightforward to calculate the spatial curvature, and it turns out to be

$${}^{(3)}R = \frac{4}{a^2}\nabla^2\varphi. \tag{3}$$

If, however, we change to new coordinates $\tilde{x}^0 = x^0 + \delta\tau(x^\mu)$ and $\tilde{x}^i = x^i + \delta x^i(x^\mu)$, then in this new gauge, $\tilde{\varphi} = \varphi + aH\delta\tau$.

Now suppose that we have defined a scalar-field perturbation $\delta\phi(\vec{x},t)$ in the spatially-flat slicing. Then, under a coordinate transformation $\tilde{t} = t + \delta t(t,\vec{x})$, $\delta\tilde{\phi} = \delta\phi - \dot{\phi}\delta t$. During slow-roll inflation, $\dot{\phi} \to 0$, so $\delta\phi$ becomes gauge-independent in this limit. Consider now the comoving gauge. In this gauge, lines of fixed $\vec{x}$ are those in which there is no momentum density; in other words, $\vec{N} = -\dot{\phi}\nabla\phi = 0$. This implies that the time displacement needed to change between the comoving and spatially-flat slicing is $\delta t = \delta\phi/\dot{\phi} \to \infty$ as $\dot{\phi} \to 0$. Therefore, the *curvature perturbation* $\mathcal{R} \equiv -\varphi = aH\delta\tau$ (a gauge-dependent quantity) becomes infinite on comoving slices (and is zero in the spatially-flat slicing). This singularity is, of course, just a coordinate singularity; it means that comoving slicing is pretty worthless when $\dot{\phi} \to 0$.

We have now defined a gauge-dependent curvature perturbation $\mathcal{R}$ which is 0 in one gauge and infinite in another....not too useful. We do expect, however, that scalar-field perturbations will in fact yield honest-to-goodness density perturbations, or at least they should if they are to be at all useful for seeding primordial structure. We'll now see how this is done. Recall that

$${}^{(3)}R = \frac{4}{a^2}\nabla^2\varphi = -\frac{4}{a^2}\nabla^2\mathcal{R} \tag{4}$$

is the curvature perturbation on spatial hypersurfaces. Let us now focus on an individual Fourier mode $\vec{k}$ of the scalar-field perturbation $\delta\phi$. We can then consider the Fourier amplitude of the density field and the Fourier amplitude $\mathcal{R}_{\vec{k}}$ of the curvature perturbation. In what follows, we will leave out the subscript $\vec{k}$ to avoid cluttered equations. Then,

$${}^{(3)}R = -\frac{4k^2}{a^2}\varphi = \frac{4k^2}{a^2}\mathcal{R}. \tag{5}$$

Since $\mathcal{R}$ is independent of the threading, this $\mathcal{R}$ will be the same for any gauge with the same slicing. Above, we argued that in the limit $\dot{\phi} \to 0$, $\delta\phi$ is the same on any slicing that is nonsingular. We can therefore assume that the analysis that gave us $\delta\phi_{\vec{k}}$ from inflation gives us $\delta\phi$ in the spatially-flat slicing. We also know from the definition of $\mathcal{R}$ that $\mathcal{R} = H\delta t$, where δt is the time-coordinate displacement between the spatially-flat and comoving slicing. But since this is $\delta t = \delta\phi/\dot{\phi}$, that means that $\mathcal{R} = H(\delta\phi/\dot{\phi})$. The next step, which we will gloss over (and perhaps leave for a homework assignment) is to note that the solution of the combined Einstein, continuity, and Euler-Lagrange equation guarantees that $\mathcal{R}$ is constant while $k/a \ll H$. Then, the power spectrum for the curvature perturbation will be, both at the end of inflation and at horizon re-entry in matter or radiation domination,

$$P_{\mathcal{R}}(k) = \left(\frac{H}{\dot{\phi}}\right)^2 \left(\frac{H}{2\pi}\right)^2\Bigg|_{k=aH}. \tag{6}$$

During slow-roll inflation, $3H\dot{\phi} = -V'$, and $H^2 = 8\pi V/3m_{\rm Pl}^2$, and so we get our final result,

$$P_{\mathcal{R}}(k) = \frac{128\pi V^3}{m_{\rm Pl}^6 (V')^2} = \frac{8}{3\pi^2}\frac{V}{m_{\rm Pl}^2 \epsilon}, \tag{7}$$

where V and ϵ are evaluated at $k = aH$. The Differential Microwave Radiometer (DMR) experiment aboard

As discussed before, NASA's Cosmic Background Explorer (COBE; 1991–4) discovered temperature fluctuations at large angular scales, or large distance scales, in the CMB. Roughly speaking, what they found was temperature fluctuations corresponding to

$$[P_{\mathcal{R}}(k)]^{1/2} = 5 \times 10^{-5}, \tag{8}$$

at $k = a_0 H_0$, from which we infer that $V^{1/4}/\epsilon^{1/4} = 0.027\, m_{\rm Pl}/8\pi = 6.6 \times 10^{16}$ GeV. Since $\epsilon \lesssim 1$, we infer that $V^{1/4} \lesssim 6 \times 10^{16}$ GeV, assuming, of course, that all of the temperature fluctuation is due to primordial density perturbations. The bound is even stronger if there is some other source of $\Delta T/T$.

Let's now understand how those $\Delta T/T$ measurements are related to $\mathcal{R}$.

The Sachs-Wolfe plateau. Fig. 1 shows data points from measurements of the CMB power spectrum $l(l+1)C_l^{\rm TT}/(2\pi)$ with the best-fit theory curve superimposed. The first order of business will be to understand the "Sachs-Wolfe plateau," the nearly flat behavior of $l(l+1)C_l^{\rm TT}$, the temperature power spectrum, for adiabatic primordial density perturbations with a Peebles-Harrison-Zeldovich (i.e., $n_s = 1$, where n_s is the scalar spectral index). More precisely, we will fill in the steps that take us from the measured temperature power spectrum $C_l^{\rm TT}$ to the amplitude $[P_{\mathcal{R}}(k)]^{1/2} \simeq 5 \times 10^{-5}$, claimed in Week 5, for the power spectrum for the curvature perturbation $\mathcal{R}$. (The "Sachs-Wolfe effect" is the term used loosely for the production of temperature fluctuations in a perturbed FRW spacetime.) Recalling that

$$P_{\mathcal{R}}(k) = \frac{8}{3}\frac{V}{m_{\rm Pl}^4 \epsilon}, \tag{9}$$

we will be able to see how measurements of $C_l^{\rm TT}$ constrain this combination of the inflation parameters.

Inflation predicts that primordial perturbations are adiabatic; i.e., that the ratios of the energy densities in each component (baryons, neutrinos, dark matter, photons) is the same from one point in the Universe to another. Its only the *total* density that varies from one point to another. What this means is that the statement that different places in the Universe have slightly different densities is equivalent to the statement that different places in the Universe are at slightly different cosmological times. In other words, cosmological events like big-bang nucleosynthesis or recombination proceed in the same way everywhere in the Universe, but at slightly different times in denser regions than in lower-density regions.

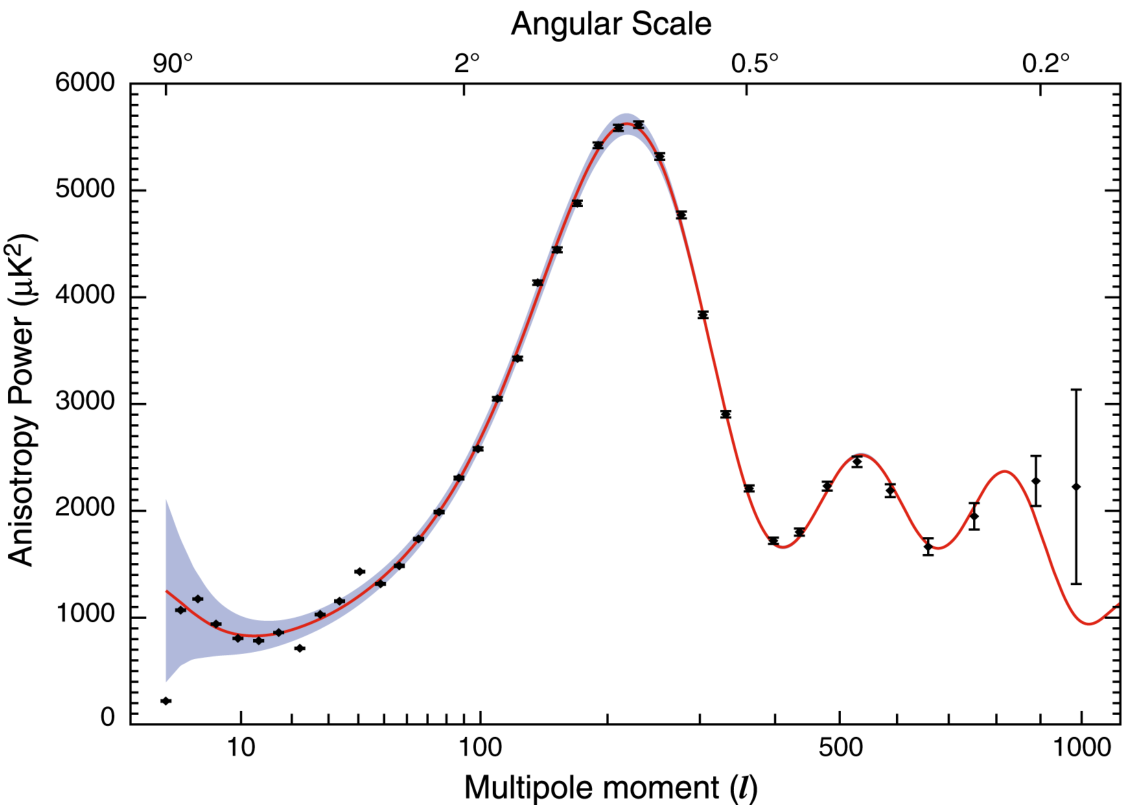

Figure 1: The temperature power spectrum $[l(l+1)/2\pi]C_l^{\mathrm{TT}}$.

Let's now consider the CMB fractional temperature fluctuation $\Delta T(\hat{n})/T_0$ as a function of position on the sky. We will consider temperature fluctuations only on large angular scales ($\theta \gg 1^\circ$), those that probe regions that are causally disconnected at the surface of last scatter. This means that the temperature fluctuations will depend only on the behavior of superhorizon modes; i.e, the density-perturbation modes that will contribute to these fluctuations have physical wavenumbers $k_{\rm phys} H \gg 1$, or physical wavelengths $\lambda_{\rm phys} \gg H^{-1}$. Thus, the amplitudes of these curvature modes will be their primordial amplitudes, with the power spectrum given above.

If the conformal time is η [defined by $d\eta = dt/a(t)$, where t is the time and $a(t)$ the scale factor], then the unperturbed metric is $ds^2 = a^2(\eta)(d\eta^2 - d\vec{x}^2)$. The distance to the surface of last scatter is thus $\delta\eta = \eta_0 - \eta_{ls}$, where η_0 and η_{ls} are the conformal time today and at last scatter, respectively.

Inflation makes a prediction for the primordial amplitude φ of the scalar metric perturbation. We now need to relate the temperature fluctuation $\Delta T/T_0$ to this metric perturbation. When we look at a direction $\hat{n}$, we are seeing the surface of last scatter a distance $\Delta\eta$ in that direction. The first thing to note is that there are two ways that a primordial density perturbation will affect the temperature we see in a given direction $\hat{n}$ on the sky. If that region is high density, then the intrinsic photon temperature will be higher. But a higher-density region corresponds to a deeper gravitational-potential well, and so the photon temperature will be redshifted as the photons climb out of that potential well. These two effects (the intrinsic temperature fluctuation and the redshift) will partially cancel. The degree to which they cancel is determined as follows.

The temperature fluctuation in a given direction is

$$\frac{\Delta T(\hat{n})}{T_0} = \frac{\Delta T_i}{T_0} - \varphi_i, \tag{10}$$

where $(\Delta T_i/T_0)$ is the intrinsic temperature fluctuation at the surface of last scatter, and φ_i is the gravitational potential in the Newtonian gauge; i.e., in the gauge where $ds^2 = a^2(\eta)[-(1-2\psi)d\eta^2 + (1-2\varphi)d\vec{x}^2]$. To simplify, we will approximate the Universe as being completely matter dominated at the time of decoupling. If so, then $\varphi = \psi$ at decoupling and afterwards. Recall also that φ reduces on sub-horizon scales to the Newtonian potential. We now need to derive a relation between $\Delta T_i/T_0$ and φ_i. Do to so, let's move to the comoving gauge, that in which time slicings are isodensity curves. In this gauge, there is no intrinsic temperature fluctuation: $(\Delta T_i/T_0)_{\rm com} = 0$. So to get the temperature fluctuation in the Newtonian gauge, we make a change in the time coordinate $t \rightarrow t + dt(\vec{x}, t)$, where $\delta t(\vec{x}, t)$ is the time shift required to get from the comoving slicing to the Newtonian gauge. Remember that aT =constant as the Universe expands, that $a \propto t^{2/3}$ during MD, and that in the Newtonian gauge, a proper-time interval is $ds = (1-\varphi)dt$. Then, under the time-coordinate shift δt required to change from comoving to Newtonian gauge, we find that in the new gauge,

$$\frac{\Delta T_i}{T_0} = -\frac{\delta a}{a} = \frac{2}{3}\varphi_i. \tag{11}$$

Putting it all together, we find the famous result,

$$\frac{\Delta T(\hat{n})}{T_0} = \frac{1}{3}\varphi(\hat{n}\Delta\eta), \tag{12}$$

relating the observed temperature fluctuation in a direction $\hat{n}$ to the Newtonian-gauge potential $\varphi(\hat{n}\Delta\eta)$ at the point on the surface of last scatter in that direction.

Now comes the calculation of the temperature power spectrum. Let's calculate the contribution of a single Fourier mode $\vec{k}$ of the density field to the temperature power spectrum $C_l^{\rm TT}$. To simplify, we'll take that Fourier mode to be in the **z** direction; since C_l's are rotationally invariant, we'd get the same result for any other direction. The temperature fluctuation induced by this forward mode, as a function of $\hat{n} = (\theta, \phi)$, is

$$\frac{\Delta T(\hat{n})}{T} = \frac{1}{3}\varphi_{\vec{k}} e^{ik\Delta\eta\cos\theta}, \tag{13}$$

where $\varphi_{\vec{k}}$ is the amplitude of the φ Fourier mode. The spherical-harmonic coefficients induced by this mode are

$$a_{lm}^{\vec{k}} = \int d\hat{n}\, \frac{1}{3}\varphi_{\vec{k}} e^{ik\Delta\eta\cos\theta} Y_{lm}(\hat{n}). \tag{14}$$

We then expand the plane wave as

$$e^{ik\Delta\eta\cos\theta} = \sum_{l'} (2l'+1) j_{l'}(k\Delta\eta) P_{l'}(\cos\theta), \tag{15}$$

where $P_l(\cos\theta) = [4\pi/(2l+1)]^{1/2}$ is a Legendre polynomial, and $j_l(x)$ is a spherical Bessel function. Orthonormality of the spherical harmonics then yields,

$$a_{lm}^{\vec{k}} = \frac{1}{3}\sqrt{\frac{4\pi}{2l+1}}\varphi_{\vec{k}}(2l+1) j_l(k\Delta\eta)\delta_{m0}, \tag{16}$$

where δ_{m0} is the Kronecker delta. Using $C_l = \sum_m |a_{lm}|^2/(2l+1)$, we find that the contribution of this mode to the power spectrum is

$$C_l^{{\rm TT},\vec{k}} = \frac{4\pi}{9}|\varphi_{\vec{k}}|^2 [j_l(k\Delta\eta)]^2. \tag{17}$$

We now have to sum over the expectation value of all Fourier modes $\varphi_{\vec{k}}$ to get the prediction for the power spectrum $C_l^{\rm TT}$. To do so requires that we specify our Fourier conventions, which we haven't yet done precisely. When this calculation is followed carefully, from the beginning of the inflationary calculation to the inflationary prediction for $P_\varphi(k)$, it turns out that the desired power spectrum is $P_\varphi(k) = (k^3/2\pi^2)|\varphi_{\vec{k}}|^2$, and so

$$\begin{aligned} C_l^{\rm TT} &= \int \frac{d^3k}{(2\pi)^3}|\varphi_{\vec{k}}|^2 \frac{4\pi}{9}[j_l(k\Delta\eta)]^2 \\ &= \int \frac{dk}{k} P_\varphi(k)\frac{4\pi}{9}[j_l(k\Delta\eta)]^2. \end{aligned} \tag{18}$$

For a Peebles-Harrison-Zeldovich power spectrum ($n_s = 1$) and these Fourier conventions, $P_\varphi(k)$ is constant as a function of k and can thus be taken outside the integral. We then use the fact that $\int_0^\infty (dx/x)[j_l(x)]^2 = [2l(l+1)]^{-1}$ to obtain

$$l(l+1)C_l^{\rm TT} = \frac{2\pi}{9}P_\varphi(k). \tag{19}$$

This shows that for a scale-invariant primordial spectrum of density perturbations (and for a matter dominated Universe), $l(l+1)C_l^{\rm TT}$ is constant on superhorizon scales $l \ll 100$, as indicated by the results of numerical CMB calculations.

The only remaining step is to figure out how $P_\varphi(k)$ is related to $P_\mathcal{R}(k)$, which is what inflation gives us. We may be at first tempted to think that $\varphi = \mathcal{R}$, as this is what we inferred was the case at the end of inflation. However, the relation $\mathcal{R} = \varphi$, does not hold at all times, and is in fact no longer true when the Universe transitions from the end of inflation to radiation domination and then later to matter domination. Instead, as shown in Chapter 6 in Dodelson's book, $\mathcal{R} = (5/3)\varphi$ for modes that enter the Universe during matter domination. (Actually, he shows a factor of 3/2, but this is the amplitude for modes that enter the horizon during radiation domination. The value of φ for modes that are still superhorizon at matter-radiation equality is decreased by 9/10 at matter-radiation equality, also as described in Dodelson's book.) We thus obtain

$$l(l+1)C_l^{\rm TT} \simeq \frac{16\pi}{75}\frac{V}{\epsilon m_{\rm Pl}^4}, \tag{20}$$

in a critical-density Universe, for $l \ll 200$. The measured value is $l(l+1)C_l \simeq 7 \times 10^{-10}$, and from this we infer the value $(V/\epsilon)^{1/4} \simeq 6.6 \times 10^{16}$ GeV mentioned in Week 5. If we want to be a bit more careful, we need to note that this calculation assumes a critical-density Universe. For a cosmological-constant Universe with $\Omega_m = 0.3$ and $\Omega_\Lambda = 0.7$, there is a small additional contribution (the "integrated Sachs-Wolfe effect") due to the growth of the gravitational-potential amplitude at late times, when the Universe becomes Λ-dominated, but this effect is relatively small, and the inflaton-potential amplitude inferred above is not changed by any more than about 10%.

The reionization optical depth. Through calculations that will be discussed the following two weeks, the first stars and galaxies to form in the Universe probably do so at a redshift $z \sim 10$. The details of this process are not well understood and it is a subject of very active current investigation. These stars emit a copious number of ionizing (energies $E \gtrsim 13.6$ eV) photons, and these knock electrons out of hydrogen atoms. With enough such ionizing photons, all of the hydrogen in the Universe becomes reionized (the "re-" occurs since the stuff was originally ionized at redshifts $z \gtrsim 1100$). Our goal here will be to calculate the optical depth τ for a CMB photon to undergo Thomson scattering sometime after recombination at redshift $z \simeq 1100$ and today. One might worry (and indeed, we once did) that if that optical depth were comparable to unity, the reionized intergalactic medium (IGM) would serve as a translucent screen, between us and the surface of last scatter, that would obscure the structure at the surface of last scatter. In fact, one finds that if the optical depth is τ, then the probability for a CMB photon from the surface of last scatter to reach us unimpeded becomes $e^{-\tau}$. Thus, the CMB power spectrum (which is quadratic in the CMB intensity fluctuation) becomes reduced by $e^{-2\tau}$ on scales smaller than the horizon at the reionization redshift. If the reionization redshift is $z_r \simeq 10$, then the corresponding multipole moments are $l \gtrsim l_r$, with $l_r \simeq 200(z_r/1100)^{1/2}/\sqrt{3} \simeq 10$ (where the 200 is the multipole moment corresponding to the sound horizon at recombination, $z \simeq 1100$, and the $\sqrt{3}$ is to change from a sound horizon to a light horizon). Fortunately, a variety of arguments now suggest that $\tau \simeq 0.1$, and so the primordial structure in the CMB is visible to us.

Before moving on to calculate τ, it is important and easy to see that only a very tiny fraction of the baryons in the Universe need to collapse into stars before reionizing the entire Universe. The basic idea is that roughly 10% of the primordial hydrogen that goes into a star gets burned to helium releasing, roughly speaking, about 8 MeV per nucleon. Approximating the hydrogen binding energy by 10 eV, we estimate that roughly 10^5 atomic binding energies worth of photons are released for each baryon processed through a star. Not all of this energy is necessarily released in ionizing photons, but stellar models suggest that a good fraction is. The actual efficiency of ionizing photons is a bit lower than one atom per 13.6 eV of photon energy, as some atoms recombine. But still, the bottom line is that only a tiny fraction of primordial gas is processed through stars before all of the IGM is ionized.

Let's assume that all of the electrons in the Universe suddenly become ionized at a redshift z_r. The optical depth to Thomson scattering is thus

$$\tau = \int_0^{z_r} n_e \sigma_T c \, (dt/dz), \tag{21}$$

where $n_e = n_{e,0}(1+z)^3$, $n_{e,0} = (7/8)\Omega_b \rho_c / m_p$ (assuming that 75% of the gas in the Universe is hydrogen and 25% is helium), Ω_b is the baryon density in units of the critical density $\rho_c = 3H_0^2/(8\pi G)$, $H_0 \simeq 73.2$ km/sec/Mpc is the Hubble constant, $m_p = 1.67 \times 10^{-24}$ g is the proton mass, c is the speed of light, dt/dz is the time interval per given redshift interval, $\sigma_T = 6.65 \times 10^{-25}$ cm^2 is the Thomson cross section, and $G = 6.67 \times 10^{-8}$ cm^3/g/sec^2 is Newton's constant. Using

$$\frac{dt/dz}{1+z} = \dot{a}/a = H(z) = H_0 \, E(z), \tag{22}$$

where (from the Friedmann equation)

$$E(z) = \sqrt{\Omega_m (1+z)^3 + \Omega_\Lambda + (1 - \Omega_m - \Omega_\Lambda)(1+z)^2}, \tag{23}$$

we find

$$\tau = \frac{3(7/8)\Omega_b \sigma_T H_0 c}{8\pi G m_p} \int_0^{z_r} \frac{(1+z)^2}{E(z)} dz. \tag{24}$$

The integral cannot be evaluated analytically easily for generaly cosmology. However, seeing that $E(z) \to 1$ as $z \to 0$ and that $E(z) \to \Omega_m^{1/2}(1+z)^{3/2}$ for $z \gg 1$, it is apparent that the integral is dominated by the large-z behavior of the integrand. We thus approximate $E(z) \simeq \Omega_m^{1/2}(1+z)^{3/2}$ to obtain

$$\tau \simeq 0.80 \frac{(z_r/10)^{3/2}(\Omega_b/0.416)(h/0.732)}{(\Omega_m/0.24)^{1/2}}. \tag{25}$$

The current best value for τ is roughly 0.1, although with sizeable errors, consistent with $z_r \sim 10$. The value is obtained by fitting the measured CMB power spectra to several cosmological parameters simultaneously, but roughly speaking, the current value comes from the amplitude of the large-angle (low-l) reionization bump in the GG (EE) polarization power spectrum, and also by comparing the amplitude of the temperature power spectrum at l's corresponding to the acoustic peak with the Sachs-Wolfe plateau at large angles (small l).

Photon diffusion and small-scale damping of the CMB power spectrum. On small scales, the CMB temperature power spectrum is damped, or equivalently, it decays with increasing l. This is due to the effects of photon diffusion or the finite thickness of the surface of last scatter, and it is sometimes referred to as "Silk damping." In words, what happens is that during the time of recombination, the photon mean-free path is increasing, until it last scatters. Consider now the penultimate scatter. Before the photon last scatters, it scatters from some other electron, and there will be a typical distance—the photon diffusion length—between the second-to-last scatter and the alst scatter. As a result, the CMB intensity we see at a given point on the sky actually samples a region at the surface of last scatter of size comparable to the diffusion length. Making the definition and calculation of the photon diffusion length precise is complicated, but roughly speaking, it corresponds to the mean thickness of the surface of last scatter.

A very rough estimate (actually, a conservative upper limit) of the photon diffusion length can be obtained in the following way. We recall that recombination takes place at a redshift $z \simeq 1100$, but that it is not instantaneous. The Universe goes from being completely ionized to almost completely neutral over a fractional redshift range $\Delta z/z \simeq 0.1$. When the Universe is completely ionized, the photon mean free path is extremely small, but when it is completely neutral, the photon mean-free path is extremely large. Thus, the diffusion length must be given by the distance corresponding to a fraction $\Delta z/z$ of the horizon at the surface of last scatter. This gives rise to a characteristic angular scale roughly one order of magnitude higher than that of the horizon. Thus, the angular power spectrum should decrease at values of l roughly $(z/\Delta z)/\sqrt{3}$ times that of the first acoustic peak, or $l \sim 1000$, which is what we see in the numerical calculations.