Missing Baryons

Michael E. Anderson

1. Outline

- Motivation

- CMB census
- Lyman-α forest census
- low-z census

- the WHIM, the CGM, and two formulations of the missing baryons problem
- How to detect the missing baryons?

suggested reading:
- Bregman, J. 2007, ARA&A 45, 221. The Search for the Missing Baryons at Low Redshift
- Fukugita, M., Hogan, C. J., and Peebles, P. J. E. 1998, ApJ, 503, 518. The Cosmic Baryon Budget
- Mo, H., van den Bosch, F., and White, S. D. M. 2010, textbook. Galaxy Formation and Evolution
- Rauch, M.1998, ARA&A 36, 267. The Lyman Alpha Forest in the Spectra of QSOs
- Shull, J. M. et al. 2012, ApJ, 759, 23. The Baryon Census in a Multiphase IGM: 30% of the Baryons May Still be Missing

2. Motivation

So far as we know, the major components of the Universe are baryons[1], radiation/relativistic particles (neutrinos and photons), dark matter, and dark energy. In the course of studying the Universe, it is natural to conduct a census of each of these components. Analytic theory and N-body simulations make clear predictions for the dark matter, and the radiation backgrounds are major fields of study as well. It turns out the majority of the photons are CMB photons, but there is also interesting science related to the infrared and X-ray backgrounds.

The baryon census is also relatively straightforward, but has led to an issue of "missing baryons" at low redshift. In this lecture I will explain how the baryon census is performed and how we know there are baryons missing, and then I will describe the reservoirs which are thought to contain these missing baryons. I will then discuss the observational attempts to detect and study missing baryons. I will show that reason the "missing baryons" are "missing" is galaxy formation,

[1]In the Standard Model of physics, baryons are composite particles made of three quarks, the most common of which are protons (Up-Up-Down) and neutrons (Up-Down-Down). Electrons are leptons, not baryons, but common nomenclature in astronomy groups them together under this category, and I will do so here as well.

and argue that studying missing baryons is a valuable way to constrain the accretion and feedback processes which are fundamental to galaxy formation.

3. Baryon Census 1: CMB and BBN

The baryon content of the Universe can be expressed either as Ω_b, the energy density of baryons, or as the baryon fraction $f_b \equiv \Omega_b/(\Omega_c + \Omega_b)$. Both of these expressions are clearly related to $\Omega_b h^2$ and $\Omega_c h^2$, which are two of the six independent parameters in the inflationary ΛCDM model of cosmology. Thus, the baryon content of the Universe is not an *a priori* prediction of ΛCDM, but must be inferred from observations. I generally prefer f_b, since it does not evolve with redshift, but I will use both in these notes. (Incidentally, in addition to its redshift evolution, Ω_b is decreasing with time ever so slightly as nuclear fusion in the cores of stars turns some of the rest mass of the baryons into photons and neutrinos. But this effect is insignificant.)

The cosmological parameters needed to compute Ω_b and/or f_b can be measured by fitting the CMB power spectrum, especially the second peak. My focus here is not on the method of CMB fitting, so I will just quote the results. The most recent Planck cosmological parameters have $\Omega_b h^2 = 0.02230 \pm 0.00014$ and $\Omega_c h^2 = 0.1188 \pm 0.0010$. Dividing these two immediately gives us $f_b = 0.158$. We can also use the Planck + external data constraints on the Hubble constant ($h = 0.6774 \pm 0.0046$) to get $\Omega_b = 0.0486$. Thus, at $z = 0$, one part in 20.5 of the Universe is baryonic, and one part in 6.3 of the matter in the Universe is baryonic.

From h, we can get the critical density $\rho_c = 3H^2/8\pi G \approx 8.6 \times 10^{-30}$ g cm^{-3}, so the mean density of ions in the local Universe is $\Omega_b \rho_c/\mu_I m_p = 1.9 \times 10^{-7}$ cm^{-3} and the mean density of electrons is $\Omega_b \rho_c/\mu_e m_p = 2.1 \times 10^{-7}$ cm^{-3} (using $\mu_I = 1.3$ and $\mu_e = 1.2$). The baryon density evolves with redshift as $(1+z)^3$.

An independent measurement can be obtained from the observed abundances of light elements (H, He, Li, Be, and their isotopes), which constrain the baryon fraction a few minutes after the Big Bang through the Big Bang Nucleosynthesis model. The result from BBN analysis is currently $\Omega_b h^2 = 0.021 - 0.025$ (see, e.g. the review by Fields et al. 2014). This is perfectly consistent with the Planck measurements at recombination ($z \sim 1100$), but much less precise.

4. Baryon Census 2: the Lyman-α forest and the Gunn-Peterson effect

After reionization, by definition most of the baryons in the Universe are ionized, and our understanding strongly points to photoionization as the dominant process. Even though the remaining neutral fraction is tiny ($\sim 10^{-4} - 10^{-5}$), it is observationally very important since it produces the Ly-α forest.

The basic physics of the Lyman-α forest is straightforward. The ionization energy of Hydrogen

is 13.6 eV, corresponding to a wavelength of 912Å for the Lyman break. The ionization cross-section for Hydrogen is huge, and although it drops off with energy as E^{-3} above 13.6 eV, neutral Hydrogen generally remains the dominant source of absorption all the way into the soft X-rays (up to ~ 1000 eV).

Below 13.6 eV, we move from bound-free to bound-bound transitions, so we have absorption lines instead of an absorption continuum. The dominant absorption line is Lyman-α, the $n = 2$ to $n = 1$ transition, which has $\lambda_{12} = 1215.7$Å. The remainder of the Lyman series is sandwiched between Lyman-α and the Lyman break: Lyman-β is the $n = 3$ to $n = 1$ transition at 1026Å, Lyman-γ is the $n = 4$ to $n = 1$ transition at 972Å, etc.

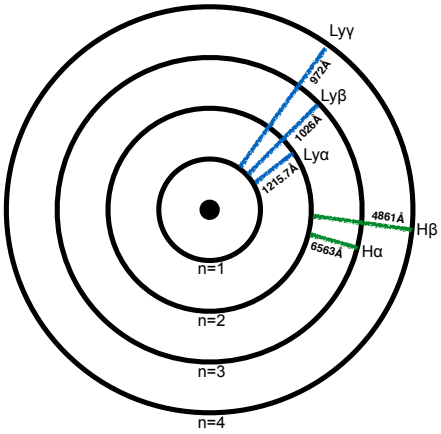

Fig. 1.— Simple schematic of the $n = 1$, $n = 2$, $n = 3$, and $n = 4$ energy levels of the Hydrogen atom, along with the first few major transitions in the Lyman (blue) and Balmer (green) series. The other important number to keep in mind is the Lyman break at 912Å, corresponding to the ionization potential of 13.6 eV for Hydrogen.

Consider a quasar at some high redshift z_Q. Neutral gas along the sightline to the quasar absorbs some its light. Let us only consider Lyman-α absorption for now. The standard equation for the absorption cross section is

$$\sigma(\nu) = \frac{\pi e^2}{m_e c} f_{12} \phi(\nu) \tag{1}$$

where f_{12} is the oscillator strength of the transition ($f_{12} = 0.416$ for Lyman-α) and ϕ is the line profile normalized such that $\int \phi(\nu) d\nu = 1$. This equation neglects stimulated emission, so it should not be applied if we have a population inversion (i.e. more neutral gas in the $n = 2$ state than the $n = 1$ state). However, the average time for a photon in the $n = 2$ state to decay to $n = 1$ is a few nanoseconds, so in the IGM we are very unlikely to have a population inversion and we can neglect stimulated emission.

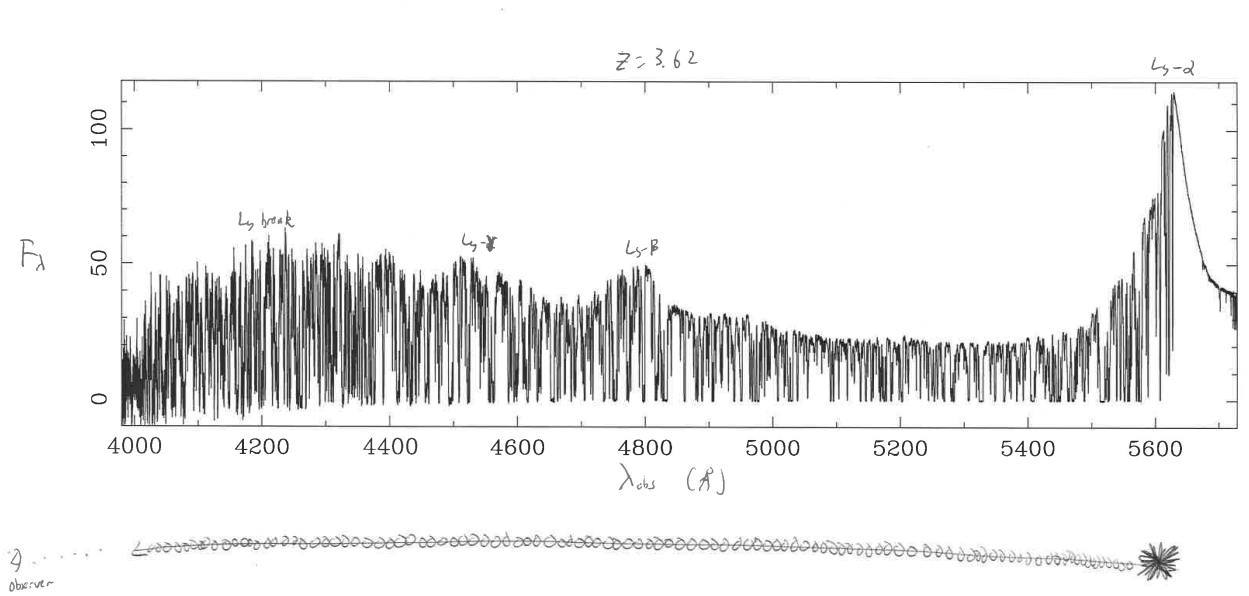

Fig. 2.— Example of the Lyman-α forest, based on a Keck HIRES spectrum of a $z = 3.62$ quasar described in Womble et al. (1996) and reproduced in Rauch (1998). I have marked the locations of some of the prominent Lyman emission lines, as well as the location of the Lyman break from the quasar. Below, I sketch a simple schematic showing the quasar sightline, and emphasizing that each absorption line in the Lyman-α forest corresponds to a single cloud containing neutral Hydrogen. In principle this forest should extend all the way to 1215.7Å (observers frame), but in practice we are limited by attenuation beyond the Lyman break, as well as confusion with other lines (Lyman lines once we are beyond Lyman-β, as well as metal lines).

Some of the FUV continuum from the quasar emitted at $\lambda < 1215.7$Å will be absorbed by intervening neutral gas as it is redshifted to the energy of Lyman-α. This is what produces the Lyman-α forest, and the forest therefore traces the column density of neutral gas in the Universe. In general, optical depth is the integral over path length of $n*\sigma$. In this case, however, σ is frequency-dependent, and over cosmological distances the proper path length is connected to redshift. This adds a layer of conceptual difficulty, but with a little thought we can still write down the optical depth from Lyman-α absorption at an observed frequency ν corresponding to somewhere in the FUV continuum of the quasar:

$$\tau(\nu) = \int_0^{z_Q} n_{\rm HI}(z) \times \sigma(\nu(1+z)) \times (dl/dz)dz \tag{2}$$

where l is the proper distance and n_{HI} is the proper density of neutral Hydrogen. Now we plug in equation 1, and approximate the line profile with a delta function. We can do this because the line profile is intrinsically narrow compared to the size of the resolution dz we achieve in practice; we are assuming that neutral gas at a single redshift z' is responsible for the Lyman-α absorption at that redshift. We obtain

$$\begin{aligned} \tau(\nu) &= \frac{\pi e^2}{m_e c} f_{12} \int_0^{z_Q} n_{\rm HI}(z)\delta(\nu(1+z) - \nu_\alpha)(dl/dz)dz \\ &= \frac{\pi e^2}{m_e c} f_{12} \left[n_{\rm HI}(z) \frac{1}{\nu} \left(\frac{dl}{dz} \right) \right]_{z=z'} \end{aligned} \tag{3}$$

where $1 + z' \equiv \nu_{12}/\nu$ (and if $z' \geq z_Q$ then we have made a bad choice of wavelength). Next we will drop the z' notation and just use z, keeping in mind that we are referring here to a single redshift somewhere between z_Q and 0. We will also convert from ν to ν_{12}, gaining a factor of $1/(1+z)$. Since ν and z are now one-to-one functions of one another, we can rewrite $\tau(\nu)$ as $\tau(z)$ with these changes:

$$\tau(z) = \frac{\pi e^2}{m_e c} f_{12} \frac{1+z}{\nu_{12}} n_{\rm HI}(z) \left(\frac{dl}{dz} \right) \tag{4}$$

Now consider the (dl/dz) factor, which is just the proper distance line element. For a photon traveling through the FLRW metric, $dl = cdt$, so we can write this factor as $c(dt/dz)$, which can then be expanded into $c(da/dz)/(da/dt)$. Since $z \equiv 1/(1+a)$, the former derivative is just $-a/(1+z)$, so we get (dropping the minus sign since we are interested in the magnitude of this factor, not its direction)

$$\frac{dl}{dz} = c\frac{da/dz}{da/dt} = c\frac{a}{\dot{a}(1+z)} = \frac{c}{H(z)*(1+z)} \tag{5}$$

where we use the definition $H(z) \equiv \dot{a}/a$. Plugging equation 5 into 4, and converting ν_{12} into λ_{12}, we obtain the final expression

$$\tau(z) = \frac{\pi e^2}{m_e c} f_{12} \lambda_{12} \frac{n_{\rm HI}(z)}{H(z)} \tag{6}$$

This is known as the Gunn-Peterson relation, and was originally derived in their 1965 paper. We can now start to re-cast this relation in terms of cosmological parameters. We know $H(z) \equiv H_0 E(z)$, where $E(z) = \sqrt{\Omega_m (1+z)^3 + \Omega_k (1+z)^2 + \Omega_\Lambda}$. Nowadays we know the Universe is roughly flat ($\Omega_k = 0$), and at the redshifts where most of the historical work has been done with the Lyman-α forest ($6 \gtrsim z \gtrsim 2$), the Universe is matter-dominated, so while I retain the $E(z)$ expression in these notes, if we need to plug in for it we can assume $E(z) \approx (1+z)^{3/2}$.

Now we introduce the connection to missing baryons. At $z = 0$, the baryon density of the Universe Ω_b is related to the Hydrogen density $n_H(0)$ through (note that n_H includes neutral and ionized Hydrogen, unlike $n_{\rm HI}$)

$$\Omega_b = \frac{n_H(0) * \mu_I m_p}{\rho_c} \tag{7}$$

where μ_I is the mean molecular weight for ions (about 1.3) and ρ_c is the critical density, $3H_0^2/8\pi G$ (roughly 8.6×10^{-30} g cm^{-3} for the Planck cosmological parameters). The neutral Hydrogen density at redshift z is related to the present-day Hydrogen density by

$$n_{\rm HI}(z) = n_H(0) * (1+z)^3 * f_{\rm HI} \tag{8}$$

where $f_{\rm HI}$ is the neutral fraction of Hydrogen at redshift z. Combining 4, 7, and 8, we obtain

$$\begin{aligned} \tau(z) &= \frac{\pi e^2}{m_e c} f_{12} \lambda_{12} \frac{3 H_0 \Omega_b}{8\pi G} \frac{f_{\rm HI}(1+z)^3}{\mu_I m_p E(z)} \\ &\approx 130 h_{70} \left(\frac{\Omega_b}{0.05} \right) \frac{f_{\rm HI}(1+z)^3}{E(z)} \end{aligned} \tag{9}$$

In Gunn and Peterson (1965), this relationship was used to constrain $f_{\rm HI}$, since the fact that we can see the FUV continuum from high-z quasars means that $f_{\rm HI}$ must be very small (i.e. the IGM is highly ionized). However, as our knowledge of the metagalactic radiation field improved, this formula has also been applied to put a lower limit on Ω_b. Such an argument is made nicely in Weinberg et al. (1997), for example, which I roughly reproduce here.

We return to equation 8, and apply our knowledge from cosmological simulations that the IGM is photoionized by the metagalactic background. Instead of anchoring $n_{\rm HI}(z)$ to the local Universe

and parameterizing the neutral fraction, we instead reformulate it in terms of photoionization equilibrium. We want the photoionization rate to balance the recombination rate. The former is just $\Gamma n_{\rm HI}$, and the latter is $\alpha n_H n_e$ (a two-body process). Γ is the photoionization rate and α is the recombination coefficient. We can equate these rates and solve for $n_{\rm HI}$:

$$n_{HI} = \frac{n_H n_e \alpha(T)}{\Gamma} = \frac{\mu_e n_H^2 \alpha(T)}{\mu_I \Gamma} \approx \frac{n_H^2 \alpha(T)}{\Gamma} \tag{10}$$

since $\mu_e \approx \mu_I$. It is also possible to incorporate a clumping factor here, if desired. For the typical Haardt and Madau (1996) metagalactic ionizing field, we expect temperatures around $1 - 2 \times 10^4$ K, for which $\alpha(T) \approx 5 \times 10^{-13} (T/10^4 \text{ K})^{-0.7}$ cm^3 s^{-1}. Then eq. 6 becomes

$$\begin{aligned} \tau(z) &= \frac{\pi e^2}{m_e c} f_{12} \lambda_{12} H(z)^{-1} \left(\frac{3 H_0^2 \Omega_b (1+z)^3}{8\pi G * \mu_I m_p} \right)^2 \frac{\alpha(T)}{\Gamma} \\ &\approx 5 \times 10^{-5} h_{70}^3 \left(\frac{\Omega_b}{0.05} \right)^2 \left(\frac{T}{10^4 \text{ K}} \right)^{-0.7} \left(\frac{10^{-12} \text{ s}^{-1}}{\Gamma} \right) \frac{(1+z)^6}{E(z)} \end{aligned} \tag{11}$$

With an assumption about the radiation field to specify T and Γ, we can therefore measure τ at a few redshifts and infer a lower limit on Ω_b (a lower limit because we are measuring the Cosmic density of baryons in the integralactic medium, and there are also baryons in galaxies). The final ingredient needed is therefore a measurement of τ. This is usually done by measuring the *flux decrement* D_A, which is related to τ via

$$D_A \equiv \left\langle 1 - \frac{F_{\rm obs}(\lambda)}{F_{\rm cont}(\lambda)} \right\rangle = \left\langle 1 - e^{-\tau} \right\rangle = 1 - e^{\tau_{\rm eff}} \tag{12}$$

The flux decrement is typically measured between Lyman-α and Lyman-β, since the Lyman-α forest begins at Lyman-α and begins to become contaminated by the Lyman-β forest after Lyman-β. Typically, the largest uncertainty is fixing the continuum, although another issue is metal line contamination (not every absorption line comes from Lyman-α, some are metal lines as well, although the IGM has a low metallicity and this is not a major effect).

Once we have measurements of τ and estimates of T and Γ from simulations or observations, we can estimate the baryon content of the IGM. The results are that the vast majority (typically more than 80%) of the total baryons in the high-redshift Universe lie in the IGM, with the rest condensed into stars and galaxies. Nowadays the estimates are more sophisticated – they use the column density distribution function of the Lyman-α forest as measured with high-resolution spectrographs, since this contains much more information than the effective optical depth. It is compared against simulations, which can account for density fluctuations and inhomogeneous reionization, but the basic spirit is the same as that presented here.

5. Baryon Census 3: The $z \approx 0$ census

We can see from eq. 11 that there is strong redshift evolution expected for the optical depth, so The Lyman-α forest can be significantly thinner in the local Universe. There is one $(1+z)^3$ factor from the expansion of the Universe and another from the increasing effectiveness of photoionization as the density decreases. There is an observational challenge as well: at low redshift Lyman-α is in the FUV, so it requires space-based telescopes in order to observe the forest. A baryon census for the low-z Lyman-α forest therefore came later than the high-redshift studies, but even in the early 1990s people were realizing there was a problem.

The problem is that there are far too few absorbers at low redshift. The most recent review of measurements is Shull et al. (2012), which estimates that the photoionized Lyman-α forest at $z = 0$ contains $28 \pm 11\%$ of the baryons in the local Universe. The errors are dominated by the photoionization correction, which is calibrated with simulations and applied to the observed column density distribution function.

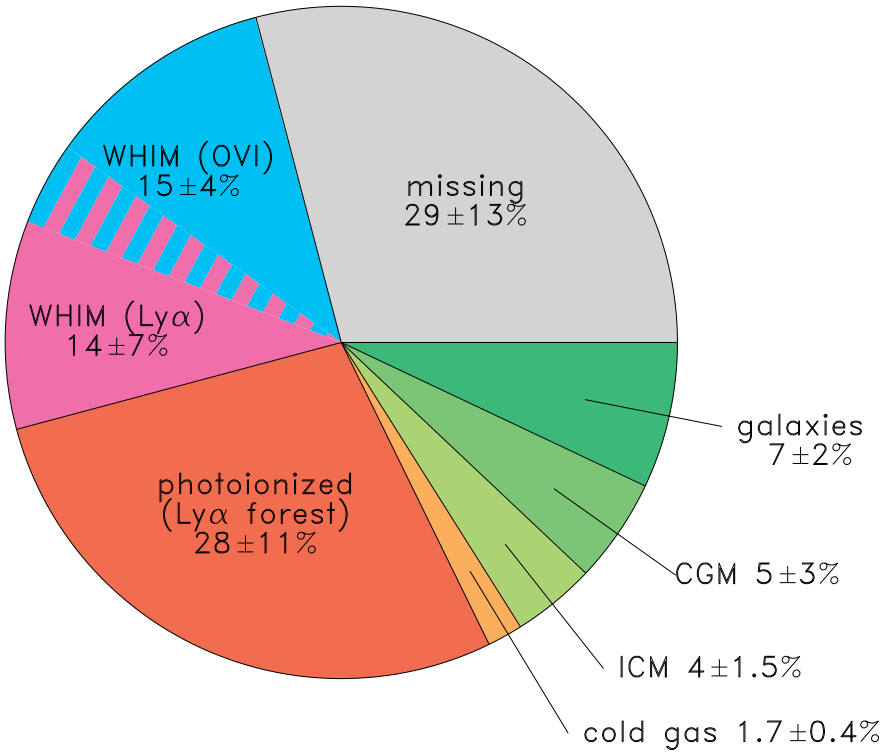

Fig. 3.— The baryon census in the local Universe, taken from Shull et al. (2012).

There are other significant baryon reservoirs in the local Universe, however. Galaxies are considerably larger today than they were 10 Gyr ago, and galaxy groups and clusters have assembled since then as well. The stars in galaxies contribute $7 \pm 2\%$ of the baryons in this paper. While IMF uncertainty may be larger than was appreciated a few years ago, clearly the stars in galaxies are subdominant. The gas in galaxies (ISM) is even less important, at $1.7 \pm 0.4\%$ of the baryons. Galaxy clusters are surprisingly important as well, at $4.0 \pm 1.5\%$ of the baryons (mostly in the intracluster medium). This is integrated down to $M_{200} = 5 \times 10^{13} M_{\odot}$, which goes down to medium-sized galaxy groups. Below this mass we can roughly switch from talking about the intragroup medium to the circumgalactic medium (CGM). Shull et al. (2012) have also estimated the mass in this phase, although the CGM is so poorly understood that even the uncertainties should be considered fairly uncertain.

So where is the remainder of the baryons in the local Universe? There is actually not much disagreement on the answer. As structures form and grow, accretion shocks heat the intergalactic medium, roughly to the virial temperature of the halo onto which the material is accreting. The virial temperature depends a bit on the density profile of the halo, but for an r^{-2} density profile with characteristic circular velocity v_c,

$$T_{\rm vir} \approx 4 \times 10^5 \text{ K} \left(\frac{v_c}{100 \text{ km/s}} \right)^2 \tag{13}$$

Comparing $T_{\rm vir}$ to the photoionization temperature of the IGM, which is around 10^4 K, we can see that accretion shocks will have a very significant impact. At high redshift, the IGM can radiate away this heat, but the cooling time scales as n_H^{-1} so as the Universe expands and the density increases, the IGM loses the ability to radiate the extra heat, and its temperature begins to increase (see

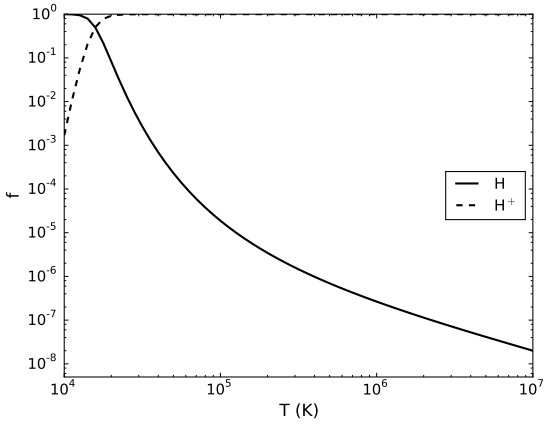

Fig. 4.— Hydrogen abundance in CIE. Note the precipitous decline of neutral Hydrogen above the typical temperatures of the photoionized IGM, a few $\times 10^4$ K.

Yannick's lecture on galaxy formation for more details). This changes the ionization equilibrium of the IGM. In the extreme case of collisional ionization equilibrium (CIE), the neutral fraction drops very quickly with temperature, as can be seen in figure 4.

Additionally, it should be noted that feedback from star formation and AGN activity in galaxies heats the IGM as well. Evidence that feedback reaches the IGM comes from measurements of the metallicity of the IGM, which seems to be roughly $0.1Z_{\odot}$ - significantly above primordial abundance. The relative importance of feedback and accretion is still an open question, but both effects push the IGM in the same direction - to hotter temperatures.

The result is that a significant portion of the IGM has been converted into the so-called *warm-hot intergalactic medium*, or the WHIM. This material no longer shows up in the Lyman-α forest, since $f_{\rm HI}$ is so low. It is thought to contain the missing baryons in the local Universe.

6. The WHIM and the CGM

In this section I want to discuss two related terms, the WHIM and the CGM, and introduce two formulations of the missing baryon problem. These are:

1. In what form are the baryons that comprise Ω_b in the local Universe?

2. For a halo of mass M, where are the $M * f_b$ of baryons which should be associated with this halo?

So far I have primarily alluded to the first version of the problem (e.g. Figure 3), but in recent years a lot of attention has also focused on the second formulation. These are not precisely equivalent to one another, since not all the matter in the Universe lies in halos, but they are obviously very related issues and I think both formulations are useful and complement one another.

To introduce the second formulation of the problem, look at Figure 5, which shows the fraction of baryons observed in halos as a function of their mass, including stars, the cold gas in the ISM, and the hot intracluster medium in galaxy clusters. While the hot phase is very uncertain (the shaded region I drew roughly covers the range of possibilities, although its extension to lower masses is not well known), it is clear that galaxies are missing a lot of their baryons, and this is the second way to look at the missing baryons problem.

The widely accepted explanation for the distinctive shape of this curve is that stellar feedback operates at the low-mass end and AGN feedback operates at the high-mass end. Since the energy available for stellar feedback scales as $M*$, while the binding energy scales roughly as $M_h^{5/3}$, stellar feedback becomes decreasingly effective at higher mass. The same scaling should apply for AGN feedback, but there is so much energy available from AGN feedback that the scaling with halo mass is not very important. A simple argument in Fabian (2012) shows this nicely. The binding energy of a bulge is roughly $E_{\rm G}{\approx}M_{\rm bulge}\sigma^2$. The central black hole mass is observed to be about $0.0014M_{\rm bulge}$

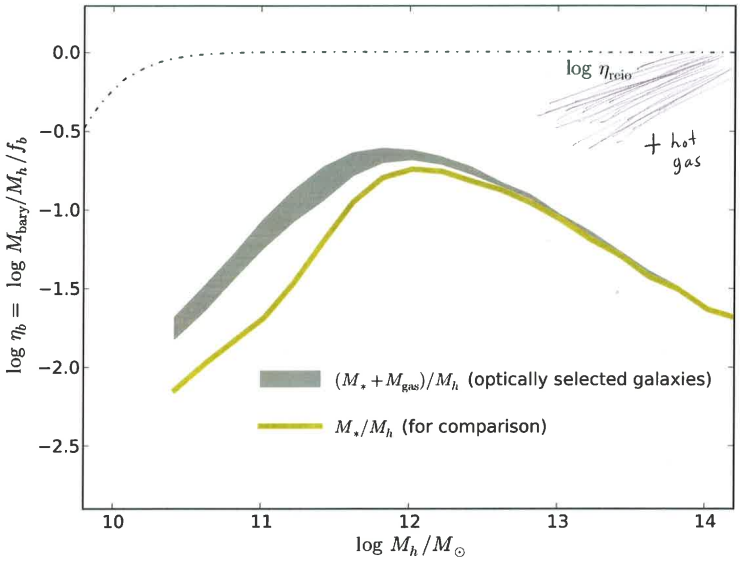

Fig. 5.— Baryon fraction of halos as a function of the halo mass, from Papastergis et al. (2012). The yellow band only accounts for stars (including subhalos), while the black band includes cold gas as traced by HI emission as well. I have shaded a region at the top right indicating the approximate range of possible baryon fractions when the intracluster/intragroup medium is included as well. The extension of this region to lower masses is still unclear. Finally, the dot-dashed line is the predicted baryon fraction without any galactic feedback except for the photoionizing background responsible for reionization.

and as it accretes it can radiate at roughly 10% efficiency, so the total energy radiated during the formation of the black hole is $E_{\rm AGN} \approx 0.1 * 0.0014 * M_{\rm bulge} * c^2$. Taking the ratio of these energies, we get

$$\frac{E_{\rm AGN}}{E_{\rm G}} \approx 140 \left(\frac{300 \text{ km/s}}{\sigma} \right)^2 \tag{14}$$

so for any reasonably sized galaxy the AGN produced more than enough energy to unbind the bulge. Thus, what matters more is the coupling of the AGN feedback to the baryons, and this coupling is more efficient at higher temperatures (because the AGN feedback is also hot).

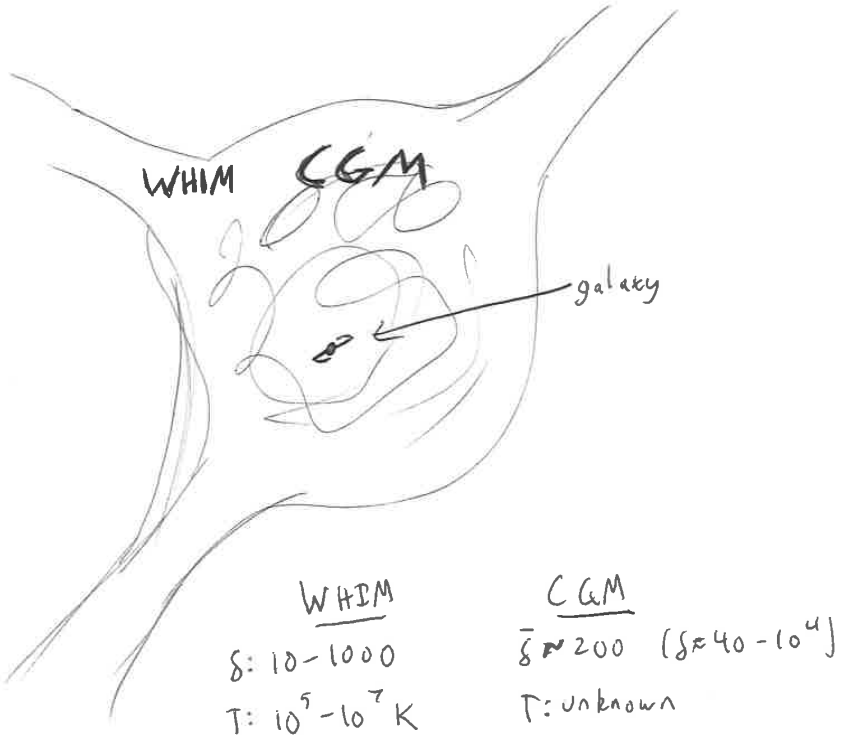

Fig. 6.— Schematic of the WHIM and the CGM, showing that these regions overlap are are not cleanly defined. For more details, see the discussion in section 6.

A useful framework is to try to determine whether the missing baryons from galaxies (1) never fell into the galaxies, (2) were ejected from the galaxies, or (3) are trapped in a reservoir around the galaxies which is known as the circumgalactic medium (CGM). The answer may be different for each galaxy, depending perhaps on the galaxy mass, environment, and/or morphology. But if one thinks of a galaxy as a machine that converts the IGM into stars, with feedback as the waste

product, then we can see that this framework is fundamental for understanding how galaxies work. The first possibility implies some sort of pre-heating at high redshifts, adding enough entropy to the IGM to prevent it from accreting onto the protogalaxies. The second possibility means that feedback is extremely powerful, both in terms of energy and momentum, and that the mass-loading factor must be significantly above unity. The third possibility should be directly observable, and has motivated dedicated searches for massive circumgalactic reservoirs.

In the rest of this lecture, I will discuss various strategies related to detecting the WHIM and its cousin, the CGM. These terms are somewhat, but not entirely, interchangeable, and I want to define them here. We can take the WHIM to be the baryons that lie in regions with overdensities in the range $10 \lesssim \delta \lesssim 1000$ and have temperatures in the range 10^5 K $\lesssim T \lesssim 10^7$ K. This roughly corresponds to large filaments and the outskirts of galaxies. The CGM is less well-defined, but it usually includes the material beyond the ISM but within either 300 kpc, $R_{\rm vir}$, or $2R_{\rm vir}$ of a central galaxy. For an NFW profile, this approximately works out to overdensities of $40 \lesssim \delta \lesssim 10^4$. The temperature structure of the CGM is not very clear yet, but the same range as the WHIM is probably reasonable.

The WHIM is thought to be fairly close in structure to the IGM at higher redshift, just hotter, more diffuse, and partially collisionally ionized. The CGM is more complicated, and its phase structure is largely unknown, but it is likely closer to a diffuse halo than a set of clouds. Thus some of the formalism developed above might not be appropriate for the CGM (i.e. the assumption of a uniform density for the absorber, and the use of a delta-function for the line profile).

7. How to Observe the WHIM

There has been an enormous amount of effort in the past 15 years towards detecting and studying the properties of the WHIM and/or the CGM. This field is moving very quickly, and at present there is roughly a paper every day or two in this vein. I can not possibly cover the entire field here, but I will try to sketch out some of the major approaches and techniques that are being used. In no particular order, these are:

- Metal absorption lines
- Diffuse emission
- Sunyaev-Zel'dovich effect (thermal and kinetic)
- other constraints

This is primarily a theory talk, so I will only sketch the basics of these methods and not go too much into a critical or comparative analysis. For those interested in more observational details, Bregman (2007) and Shull et al. (2012) perform extensive reviews, although they primarily focus on the WHIM. I am not aware of a good review of the current status of the CGM, but also I do not think there is much of a consensus yet on the observations and how to interpret them.

7.1. Metal absorption lines

I will focus here on Oxygen, whose intermediate and high-ionization states I have sketched out below. Oxygen is one of the most useful tracers of the IGM because of its high abundance and because O VI, O VII, and O VIII span nearly the entire 10^5-10^7 K temperature range covered by the WHIM (with C IV and Lyman-α contributing at the low end as well). O VI is a UV doublet with lines at 1032Å and 1038Å, which also makes it easy to identify. Once the lines are identified, the column density distribution function $F(N_{\rm OVI})$ is derived, and then, in a slightly more sophisticated version of the analysis in section 4, we have

$$\Omega_{\rm OVI} = \frac{H_0}{c}\frac{m_O}{\rho_c}\int_0^\infty F(N_{\rm OVI})\ N_{\rm OVI}\ dN_{\rm OVI} \tag{15}$$

Unfortunately, compared to Lyman-α, $\Omega_{\rm OVI}$ is rather difficult to interpret. One issue is that O VI turns out to be relatively easy to produce either from photoionzation or collisional ionization. In CIE, the maximum fraction of Oxygen which is in O VI is only ~ 0.2 (see Figure 7), so the ionization correction can be considerable, and particularly in the CGM it can be difficult to be sure whether CIE applies for O VI. An additional issue is that we need to specify a metallicity in order to convert from $\Omega_{\rm OVI}$ into Ω_b, and the metallicity is usually unknown (and is calibrated using simulations, typical values are $0.1 - 0.2Z_\odot$).

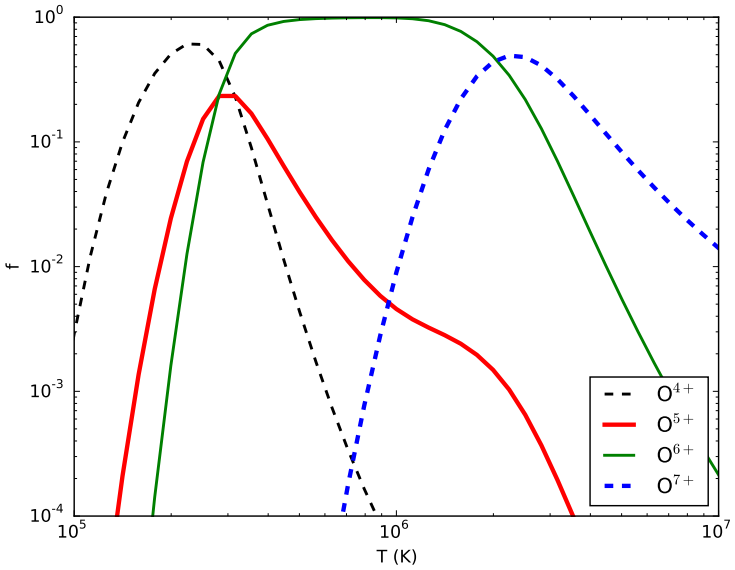

Fig. 7.— Oxygen abundance in CIE. Recall that O V lines refer to transitions in O^{4+}, with similar notation for the other states. Note also that O^{5+} is subdominant at every temperature.

Even though the neutral fraction is suppressed in the WHIM, near 10^5 K at least, it can still be traced by Lyman-α absorption as well. These Lyman-α lines look different from the narrow lines typical of the photoionized gas, however, since they are also appreciably thermally broadened. These are known as broad Lyman-α absorbers, or BLAs, and overlap with the O VI-traced WHIM in Figure 4 (the overlap is due to uncertainty about the interpretation of the O VI absorbers). The amount of broadening we expect is given by the Doppler broadening parameter $b = \sqrt{2kT/m_I}$, and for Hydrogen this works out to $b \approx 40$ km/s $*(T/10^5$ K$)^{1/2}$.

O VII and O VIII correspond to a series of soft X-ray lines around 0.6-0.8 keV. With moderate redshift applied and absorption from neutral Hydrogen, this is a difficult regime for spectral analysis, although accessible with the grating spectrographs on XMM-Newton and Chandra (useful down about to 0.35 and 0.2 keV respectively). Unfortunately, current X-ray gratings have effective areas of tens of cm^2 (compared to 1000-3000 for HST-COS in the FUV), so they have very little sensitivity to the WHIM. There have been a few claims of O VII detections of WHIM filaments, but most have not been confirmed. The few that have (e.g. the Sculptor wall) are not really sufficient to tell us much about the baryon content in the hotter phases of the WHIM.

On the other hand, O VII and O VIII are very commonly detected at zero redshift, corresponding to the hot halo around the Milky Way. Additional leverage came from the inclusion of O VII and O VIII *emission* lines, which can be measured in any deep blank field observation. O^{6+} and O^{7+} are mostly ionized collisionally, so their emissivity depends on the square of the gas density, while absorption is only linearly dependent on the density. With enough lines, plus the different density dependences, it has finally been possible to begin to decompose the hot halo from the hot interstellar medium in the Galaxy. The picture that emerges points to a hot halo containing a mass roughly similar to the stellar mass of the Galaxy - i.e. a significant amount of baryons, but not the full complement of missing baryons from the galaxy (see, e.g. Miller and Bregman 2015).

7.2. Diffuse emission

The CGM and the WHIM also produce diffuse emission which can be detected in some cases. The cooler photoionized gas emits Lyman-α and the emissivity is linearly proportional to the gas density. Hotter collisionally ionized gas emits soft X-rays, but the emissivity is proportional to the square of the gas density, which limits the effectiveness of emission as a probe of low-density plasma.

Attempts have been made to image both the Lyman-α halos and the hotter X-ray halos. The Lyman-α analysis is typically restricted to higher redshifts so the emission is shifted into optical bands where it is easier to observe, but there are impressive results coming from dedicated instruments like the Cosmic Web Imager in this direction. The X-ray analysis can only be done at low redshift, both because the emission is so intrinsically faint and because at higher redshifts it shifts out of the X-ray bandpass entirely.

Lyman-α traces Hydrogen directly, but in order to infer a gas mass, one needs to know the ionization correction and the escape fraction, both of which can be considerable corrections and are controversial.

X-ray analysis traces metal emission lines (for the CGM, primarily O VII and O VIII, but also some Iron lines), but X-ray telescopes have a major advantage in that they simultaneously measure the energy and position (and arrival time) of each photon, so the spectrum can be measured along with a surface brightness analysis. In principle, this allows us to measure, for example, the metallicity of the hot gas so the total mass can be inferred. The metallicity can be inferred by measuring the ratio of metal emission lines to the continuum, which is usually done automatically with X-ray spectral modeling software. In practice, however, this is extremely difficult. One issue is that current CCD detectors cannot resolve the metal emission lines from plasma at CGM temperatures, and another problem is that the X-ray continuum (from thermal bremsstrahlung) is intrinsically very faint and can often be overwhelmed by the emission lines themselves. Usually a reasonable choice is to guess a metallicity ($0.2Z_\odot$ and $0.5Z_\odot$ are popular choices) and then assume the mass scales from this fiducial value as $Z^{-0.5}$, but obviously this is insufficient if the goal is 10% or better precision. The general conclusion from this analysis agrees with the Milky Way results, however: around the handful of giant galaxies with measurements, hot halos as traced by X-rays seem to contain similar amounts of matter as the stars in these galaxies. This either implies that there are still additional baryons missing from the CGM, or that the remaining hot baryons are lying at larger radii where they are too diffuse to detect in emission.

In addition to direct imaging of individual objects, other methods are also used to improve the S/N. One method is stacking (or equivalently, cross-correlation) where some other tracer of large-scale structure is used along with the imaging analysis. For X-rays, this could mean stacking the emission around galaxies, or it could mean cross-correlating the X-ray background with a galaxy catalog or with the NIR background. Another possibility is generating a power spectrum of the diffuse background (or equivalently, autocorrelation). This is more challenging and historically has not yielded results as robust as stacking, but progress is still being made in this direction as well.

7.3. Sunyaev-Zel'dovich effect

The Sunyaev-Zel'dovich effect is the imprint on the CMB produced by inverse Compton scattering between CMB photons and hot electrons. Two related effects can be distinguished, the thermal SZ effect and the kinematic SZ effect. The previous cosmology lecture by Daisuke Nagai has an excellent introduction to these effects.

In the non-relativistic limit (as applies to the WHIM and CGM, which have temperatures $\lesssim$ 1 keV compared to the electron rest mass energy of 511 keV), and in the Rayleigh-Jeans portion of the spectrum (i.e. at frequencies longer than 217 Ghz), both effects are proportional to the optical depth τ of Thomson scattering from free electrons along the line of sight (integrated from Earth

to the last-scattering surface at $z \sim 1100$). Since the WHIM and the CGM are highly ionized, it is worth considering the utility of the SZ effect for studying these media.

The thermal SZ effect is a decrement in the background caused by scattering against the thermal motion of electrons of temperature T_e. Its magnitude is

$$\frac{\Delta I}{I} \approx 2\tau \left(\frac{kT_e}{m_e c^2} \right) \approx 3 \times 10^{-5} \tau \left(\frac{T_e}{10^5 \text{ K}} \right) \tag{16}$$

The kinematic SZ effect is a decrement in the background caused by scattering against the peculiar velocity v_e of the electrons relative to the CMB frame.

$$\frac{\Delta I}{I} \approx \tau \frac{v_e}{c} = 10^{-3} \tau \left(\frac{v_p}{300 \text{ km/s}} \right) \tag{17}$$

This leads to the somewhat surprising result that the kinematic SZ effect is actually generally larger than the thermal SZ effect for the WHIM and the CGM. This is the opposite of the case for galaxy clusters, where the thermal SZ effect dominates in most cases. Unfortunately, the optical depth to Thomson scattering is tiny, so neither effect is easy to observe. For example, considering a good case such as looking down a filament of length 1 Mpc and electron density 2×10^{-6} cm^{-3} (i.e. $\delta \sim 10$), the optical depth is just $\tau \sim 4 \times 10^{-5}$. The CGM is a bit more promising and can produce larger optical depths, but obviously the ICM of a massive galaxy cluster will always be bigger than the CGM of a massive galaxy.

For the CGM, the Planck Collaboration stacked the thermal SZ maps around central galaxies and was able to detect a signal around central galaxies of just over $10^{11} M_\odot$ in stellar mass. Several groups have looked into the prospects of detecting the WHIM with the SZ effect, and it generally seems to be out of the reach of current SZ surveys, but perhaps possible in the future. Cross-correlating SZ maps with other tracers of structure is another possibility.

7.4. Other constraints?

Finally, it is worth mentioning that two other parameters can also be measured which can help to constrain the intergalactic medium - the optical depth to reionization τ and the mean thermal SZ signal in the sky $\overline{Y}$. These have both been measured by Planck, and are sensitive respectively to the integrated column density of free electrons and the integrated free electron pressure in the Universe (out to $z \sim 1100$). In principle these quantities are related to the baryon content of the IGM, but as they are averaged over the whole Universe they are not especially easy to work with.

There have also been a few creative attempts to study the feasibility of measuring the dispersion measure for extragalactic radio sources, which gives a direct measurement of the free electron column density between us and these extragalactic sources. I am not aware of much progress

in this area[2], but if such measurements could be made, they would also provide very valuable constraints.

REFERENCES

Bregman, J. N. 2007, ARA&A, 45, 221

Fabian, A. C. 2012, ARA&A, 50, 455

Fields, B. D., Molaro, P., & Sarkar, S. 2014, arXiv:1412.1408

Gunn, J. E., & Peterson, B. A. 1965, ApJ, 142, 1633

Miller, M. J., & Bregman, J. N. 2015, ApJ, 800, 14

Papastergis, E., Cattaneo, A., Huang, S., Giovanelli, R., & Haynes, M. P. 2012, ApJ, 759, 138

Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2015, arXiv:1502.01589

Rauch, M. 1998, ARA&A, 36, 267

Shull, J. M., Smith, B. D., & Danforth, C. W. 2012, ApJ, 759, 23

Weinberg, D. H., Miralda-Escudé, J., Hernquist, L., & Katz, N. 1997, ApJ, 490, 564

Womble, D. S., Sargent, W. L. W., & Lyons, R. S. 1996, Cold Gas at High Redshift, 206, 249

[2]The dispersion measure has been measured for some fast radio bursts, but for these sources we do not know the distance, so constraints are still not very useful.

This preprint was prepared with the AAS LaTeX macros v5.2.