1. CMB Observations and cosmological constraints

R. BRUCE PARTRIDGE

Abstract

Measurements of the spectrum of the cosmic microwave background (CMB) and of the power spectrum of fluctuations in its intensity across the sky have allowed us to refine crucial cosmological parameters, and have opened up the age of "precision cosmology." Recent spectral and anisotropy measurements of the CMB are discussed; some of the observational difficulty faced in making such observations is noted, from an experimentalist perspective. Then some of the conclusions which can be drawn from measurements of the spectrum and of the power spectrum of fluctuations, as well as from increasingly sensitive measurements of polarized fluctuations in the CMB, are presented.

1.1 Introductory remarks

My aim in the talks given at the XIX Canary Islands Winter School of Astrophysics was to present recent observational work on the cosmic microwave background (CMB), to sketch some of the observational difficulties in making CMB measurements and to list the cosmological constraints that those measurements established. I included a good deal of pedagogical material which I have reduced, given limits of space, in this written version. The written version relies more on references to the literature and emphasizes the observational constraints that several decades of careful measurements of the CMB have established. I continue to give the material an observational or experimental flavor, since most of the other participants in the Winter School were theorists. I also retain the informal style of my talks in Tenerife.

1.2 The spectrum of the CMB and what it tells us

The story of the discovery of the cosmic microwave background radiation (CMB) is well known, and frequently recounted as an example of serendipity in science. In 1964, Arno Penzias and Bob Wilson were faced with the problem of "excess temperature" in a communications antenna operating at a wavelength of 7 cm. The signal was difficult to understand because it was substantial in amplitude (equivalent to 3.5 K) and "isotropic, unpolarized and free from seasonal variations," to quote their 1965 paper (Penzias and Wilson). Note that Penzias and Wilson had already remarked on the isotropy of what we now know as the CMB. The actual interpretation of the signal as cosmic in origin was not made by Penzias and Wilson, but in a companion paper by Dicke *et al.* (1965), who argued that it is radiation left in the Universe from an earlier high-density, high-temperature, ionized phase.

That interpretation was not immediately accepted by all astronomers. We needed to show the ∼3 K signal was cosmic, not local, in origin. There were two important tests of the cosmic hypothesis: the spectrum should be blackbody, and the radiation should be largely isotropic. The isotropy follows from the approximate isotropy and homogeneity of the Universe as a whole.

Now let us explore why a blackbody spectrum would be expected from a hot early phase. Independent of any particular value for the current baryon density, there must have been a time early in the history of the Universe when the timescale for thermal processes was shorter than the expansion timescale, H^{-1}. For reasonable values of the current baryon density, it can be shown that the corresponding redshift was $\sim 2 \times 10^6$ (Burigana *et al.*, 1991; see also chapter 5 of Partridge, 1995). Unless pathological initial conditions are assumed, the Universe would have had time to reach thermal equilibrium by a time corresponding to that redshift, or roughly one month after the Big Bang. Both scattering processes, like the Compton and inverse Compton effects, and photon-generating processes, such as thermal bremsstrahlung and the radiative Compton effect, were involved. Thermal equilibrium in turn generates a blackbody spectrum in the radiation field.

Because the expansion of the Universe is adiabatic, a blackbody spectrum, once established, is maintained, and expansion affects only the temperature as $T(z) = T_0(1 + z)$. The argument that adiabatic expansion preserves a blackbody spectrum is made in Weinberg's book on *Gravitation and Cosmology* (1972), and goes back to much earlier work by Robertson.

By 1967, observations had shown that the radiation was isotropic to a few tenths of a percent, and that its spectrum was consistent with a blackbody curve with a temperature of $\sim$2.75 K over a range of wavelengths exceeding a factor of ten (see e.g. Stokes *et al.*, 1967; Wilkinson, 1967 and a summary of measurements using interstellar cyanogen by Thaddeus, 1972). The interpretation of Dicke and his colleagues had passed both tests. In addition, of course, attempts to improve the spectrum led to a more precise value of the present temperature, T_0.

As the cosmic interpretation of the microwave excess took hold, increasing interest arose in small departures from perfect isotropy and small departures from a perfect blackbody spectrum. In this section, I explore the latter.

1.2.1 Varieties of equilibria, and the traces they leave in the CMB spectrum

I will condense here material discussed in somewhat more detail in Chapter 5 of Partridge (1995). First, *kinetic equilibrium* can be established by any scattering process with a timescale substantially less than H^{-1}. Under the conditions of kinetic equilibrium, as opposed to true thermal equilibrium, the occupation number of photons can be written as $\eta = \left[\exp\left(\frac{h\nu}{kT} + \mu\right) - 1\right]^{-1}$ where μ is the "chemical potential." The resulting spectrum is the Bose–Einstein spectrum. Clearly, μ plays only a small role at very high frequencies, where the above expression for η relaxes to the Planckian form; the discrepancy becomes larger and larger as the frequency drops.

True *thermal equilibrium* requires the creation and destruction of photons as well as energy redistribution by scattering (see, for instance, Kompaneets, 1957). In the early Universe, both radiative Compton and bremsstrahlung processes permit the generation of photons needed to ensure true thermal equilibrium. While the former is likely to be dominant with current values of the baryon density (Burigana *et al.*, 1991), the latter process generates a characteristic spectral feature worth exploring. The bremsstrahlung cross section is proportional to λ^2, and as a consequence bremsstrahlung is most effective at creating thermal equilibrium, and a consequent blackbody spectrum, at large wavelengths. Now consider the possibility that bremsstrahlung has not completed the process of thermal equilibration. The result will resemble Fig. 1.1, where the dashed line is a Bose–Einstein spectrum characterized by a positive chemical potential, μ, and the solid line shows the effect of partial equilibration by bremsstrahlung at low frequencies.

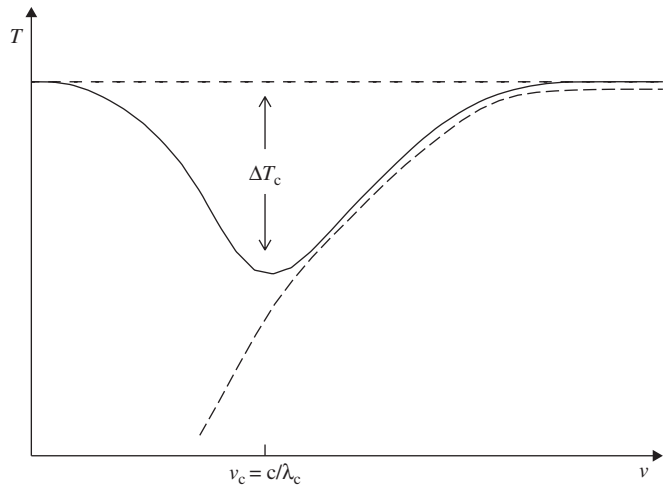

FIGURE 1.1. The "μ-distortion" of the CMB spectrum described in the text; the dashed line shows a pure Bose–Einstein spectrum.

The claim that thermal equilibrium is established at the high redshift of $\sim 2 \times 10^6$ is equivalent to the claim that μ is driven essentially to zero by that redshift. However, if energy is added to the CMB radiation field *after* an epoch corresponding to a redshift of $\sim 2 \times 10^6$, there may still be time to reintroduce kinetic equilibrium, but not full thermal equilibrium. In this case, the spectrum shown by the solid line in Fig. 1.1 will result. We refer to the departure from a Planck spectrum at fixed T as a "μ-distortion." The center frequency (or wavelength, λ_c) and the amplitude of the evident dip depend on μ and on the baryon density (since only baryonic matter is involved in the thermalizing processes). For instance, $\lambda_c \sim 2.2\left(\Omega_b h^2\right)^{-2/3}$ cm. And ΔT_c, the amplitude of the dip, is given by $\Delta T_c/T_0 = 2.3\,\mu(\Omega_b h^2)^{-2/3}$ (from Illarionov and Sunyaev, 1975; and Burigana *et al.*, 1991). Since λ_c depends only on Ω_b and h, and not μ, we can evaluate it: $\lambda_c \sim 30$ cm for currently accepted values of the cosmological parameters.

Now let us ask about scattering processes, those that conserve photon number, but readjust photon energy. By far the dominant process is inverse Compton scattering, in which the energy of a CMB photon is "boosted" by collision with hot electrons in the ionized material. This is the effect explored by Sunyaev and Zel'dovich (1969 and 1980a,b). This scattering process in a uniform plasma affects the overall spectrum of the CMB. The resulting spectral distortion can be written as

$$\frac{\delta I(\nu)}{I(\nu)} = -y \frac{x e^x}{e^x - 1}\left[4 - x \coth\left(\frac{x}{2}\right)\right],$$

derived, for instance, by Peacock (1999) from the Kompaneets equation. Here $x = \left(\frac{h\nu}{kT_0}\right)$ and the y parameter involves the integral along the line of sight of the electron pressure,

$$y = \int \frac{kT_e n_e \sigma_T}{m_e c^2} \mathrm{d}\ell. \tag{1.1}$$

In this expression T_e, n_e and m_e are the temperature, number density and mass of the electrons, respectively, and σ_T is the Thompson scattering cross section.

1.2.2 Experimental issues

To look for either form of distortion, sensitive measurements of the CMB spectrum over a substantial range of wavelength are required. Here, I discuss briefly some of the observational difficulties in making such measurements. First, for an *absolute* measurement of T_0 at a given wavelength, one needs some form of absolute calibration. While in principle a blackbody at any known temperature could be used, it is strongly advantageous to have a calibration blackbody at a temperature close to 2.7 K, so that non-linearities in the response of the detector do not bias the measurement. For that reason, all CMB spectral measurements have employed cryogenic calibrators, referred to (in terminology introduced 60 years ago by Dicke) as "cold loads." An ideal cold load would enclose a measuring antenna with a perfect blackbody at a precisely measured temperature. Ground-based measurements have used large dewars containing liquid helium and a (nominally) perfect blackbody immersed in it (e.g. Smoot *et al.*, 1985). The FIRAS instrument aboard the Cosmic Background Explorer (COBE) satellite, to be discussed below, employed a somewhat similar system, but allowed for small temperature adjustments in the cold load to equal and thus null out the CMB signal.

If an accurate calibration is achieved, the remaining problems are sources of foreground radiation that add to the CMB signal entering an antenna viewing the sky. These take three forms: radiation from the ground and nearby structures that diffracts into the antenna mouth; radiation from the Earth's atmosphere; and radiation from the Galaxy. The first of these can be mitigated by the use of ground screens, reflective surfaces that prevent direct radiation from the hot ground into the antenna mouth. Ground screens are now routinely used for *all* CMB observations, whether they are spectral or isotropy observations. Emission from the atmosphere can be measured and subtracted by using the zenith angle dependence of atmospheric emission. In the case of a thin, plane-parallel atmosphere, the atmospheric signal scales simply as $T_{\rm atm} \propto \sec z$, where z is the zenith angle. Atmospheric emission can also be reduced by observing at frequencies between the strong water-vapor and oxygen lines that dominate microwave emission of the atmosphere. The presence of the strong H_2O line at 22 GHz and O_2 lines at 60 and 120 GHz explain why so many ground-based observations are made at frequencies well below 22 GHz, or in "valleys" between the lines, that is near 30 or 90 GHz. Atmospheric emission is discussed in more detail in the review by Weiss (1980) and in Danese and Partridge (1989). Some rough numbers show the magnitude of the problem presented by the atmosphere: at sea level $T_{\rm atm} \gtrsim 3\,{\rm K}$ (or $\gtrsim T_0$) for $\lambda \lesssim 3\,{\rm cm}$. Even at the best, high altitude sites, $T_{\rm atm} \sim 2\,{\rm K}$ for $\lambda = 3\,{\rm mm}$.

Note that atmospheric emission pushes one to observe at relatively low frequency. Unfortunately, the Galactic background favors observations at high frequency. In the microwave region, the dominant source of Galactic emission is the synchrotron process, with a temperature spectrum going roughly as $T_{\rm G} \propto \nu^{-2.8}$, a proportionality that holds for frequencies below or about 90 GHz. At higher frequencies, reemission from Galactic dust becomes important, as does free–free emission with a flatter spectrum, $T_{\rm G} \propto \nu^{-2.1}$. We will return to Galactic emission in Section 1.3.3.2. Of course, Galactic foregrounds can be minimized by making spectral measurements near the Galactic poles.

1.2.3 COBE

By the mid 1970s, it was clear that improved spectral and anisotropy measurements, at least on a large scale, required a space platform. The COBE satellite, launched in 1989, was designed to provide both kinds of measurements, and to probe far infrared radiation from the Galaxy as well. The spectrum was measured by the FIRAS instrument, the "Far Infrared Absolute Spectrometer." Unlike earlier ground-based measurements, it

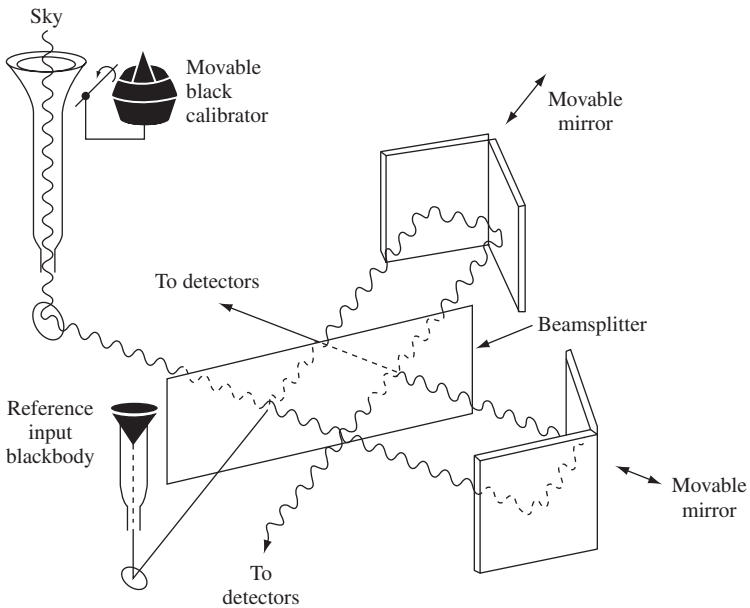

FIGURE 1.2. Block diagram of COBE's FIRAS instrument. (From COBE slide set. NASA image.)

employed a Michelson interferometer to measure temperature as a function of wavelength over a wavelength range of roughly 0.05–0.5 cm. The instrument was carefully designed to make a null measurement, by comparing the observed temperature of the CMB with the temperature-controlled cold load. As a further check, the external antenna could be covered by a second temperature-controlled cold load, as shown in the schematic in Fig. 1.2. It is worth emphasizing that uncertainties in the final result from FIRAS depend more on understanding the two cold loads than they do on actual measurements of the sky.

The results of the first few minutes of COBE's observation were presented in January of 1990, to a standing ovation (Mather *et al.*, 1990). Subsequent analysis, again particularly of the cold load, has resulted in a slight improvement in the sensitivity of the experiment, and I give the final result (Fixsen and Mather, 2002):

$$T_0 = 2.725 \pm 0.001\,\mathrm{K}. \tag{1.2}$$

At about the time the COBE results were first announced, a pioneering rocket experiment, also based on an interferometer, was measuring the CMB spectrum over a similar range of wavelengths. This important work by Gush *et al.* (1990) provided crucial confirmation of the COBE results.

The precision of the 2002 COBE result is astonishingly high – we know the temperature of the CMB, a crucial cosmological parameter, to better than 0.05%. However, the wavelength range covered by COBE was limited, particularly at the long-wavelength end. As a consequence, there have been and continue to be attempts to improve our knowledge of the CMB spectrum at wavelengths of 1 cm and longer. In particular, to reveal a possible μ-distortion, observations at $\lambda > 10$ cm are needed.

1.2.4 Longer wavelength measurements

At wavelengths longer than about 3 cm, the atmospheric emission is sufficiently small to allow these experiments to be done from the ground, though balloons are also being

used, for instance by the ARCADE group (Kogut *et al.*, 2006). To date, the most precise observations are those made by the TRIS team based in Milan. While these results have not yet been published, Giorgio Sironi has kindly given me permission to quote some numbers; for details see Gervasi *et al.* (2008).

- At 50 cm, the group found $T_0 = 2.685 \pm 0.038$ ($\pm$ 0.066) K
- At 36 cm, $T_0 = 2.772 \pm 0.012$ ($\pm$ 0.40) K
- At 12 cm, $T_0 = 2.516 \pm 0.107$ ($\pm$ 0.28) K

Here, the numbers in parenthesis include systematic error, largely due to correction for Galactic emission. To these we may add the following observations at slightly shorter wavelengths:

- At 3 cm, $T_0 = 2.730 \pm 0.014$ K (Staggs *et al.*, 1996)
- At $\sim$4 cm, $T_0 = 2.84 \pm 0.014$ K (Kogut *et al.*, 2006)

1.2.5 Constraints on spectral distortions

We have known for more than a decade (see summary in chapter 4 of Partridge, 1995) that any distortions in the CMB spectrum are small. Let us now examine the best current limits on distortions and what they imply.

***y*-distortion.** The COBE team (Fixsen *et al.*, 1996) fixed a tight upper limit on the y parameter: $y < 1.5 \times 10^{-5}$. This in turn limits the amount of inverse Compton scattering by a generally distributed plasma at $z \lesssim 10^5$ (local concentrations of plasma, such as the intergalactic medium (IGM) in clusters of galaxies, can produce much larger values of y; these are discussed in Section 1.4 below).

The refinement of the COBE data in later papers by the COBE team (e.g. Fixsen and Mather, 2002), and the addition of long-wavelength measurements allowed the TRIS team to establish improved limits (Gervasi *et al.* 2008). They claim $-5 \times 10^{-6} < y < 3.5 \times 10^{-6}$. Since there is a direct connection between y and the excess energy added to the radiation field Δu, we can in turn limit Δu:

$$\Delta u/u = 4y, \text{ so that} \qquad \Delta u \lesssim (2 \times 10^{-5})u, \tag{1.3}$$

where u is the energy density of the CMB radiation field.

One potential source of a y-distortion is a thin, more-or-less uniform, intergalactic plasma. There was a vogue some years ago for explaining a portion of the X-ray background as thermal emission from such a hot IGM, with $T \sim 10^6$ K. The observed limits on y set extremely tight limits on the density of such a plasma (this is left as an exercise; see Wright, 1994, for some guidance).

***μ*-distortion.** Here the constraint from TRIS and COBE measurements is $\mu < 7 \times 10^{-5}$. For small values of μ, $\Delta u/u \sim 0.7\,\mu$ (Illarionov and Sunyaev 1975) so we find $\Delta u \lesssim 5 \times 10^{-5} u$. This limits energy release (say, by decay of exotic particles or the evaporation of low-mass black holes by the Hawking mechanism) at epochs less than $\sim$30 y.

1.3 Measurements of CMB anisotropies on large angular scales ($\ell \lesssim 300$)

Those who attended the Winter School in Tenerife will recognize that I have changed the order of this section. This was done largely to remove some overlap between my presentation, Wayne Hu's, and to a lesser extent Sabino Matarrese's.

1.3.1 A cartoon view of a power spectrum of anisotropies

I will begin by presenting a cartoon version of the power spectrum of CMB anisotropies; for details see the much more careful presentations by Hu and Matarrese. This quick outline is designed to highlight the key physical processes responsible for features in different ranges of spatial frequency, ℓ, in the power spectrum of the CMB (Fig. 1.3). This approach allowed me to divide up the presentation of CMB anisotropies into two sections, corresponding to two talks at the Winter School. Recall that $\ell \sim 180^\circ/\theta$.

$\ell \lesssim$ 50–100. The angular scale of a causal horizon on the surface of last scattering (at $z \sim 1000$) corresponds approximately to 2 degrees. At angular scales larger than this, or $\ell \lesssim 100$, we thus see the power spectrum imprinted during the inflationary epoch, unaffected by later, causal, physical processes. This point is explained in detail in Matarrese's article here. He also explains why the shape of the power spectrum is expected to have a power-law index close to $n_S = 1$, which in turn translates into a nearly flat power spectrum at values of $\ell \lesssim 100$.

At these low values of ℓ, we also need to consider cosmic variance. This is perhaps a slight misnomer: the variance is not in the cosmos but in the models. Consider a specific cosmological model, with chosen values for the various cosmological parameters. The model parameters can be used, for instance in programs like CMBFAST, to calculate the CMB power spectrum. At low values of ℓ, however, there is some intrinsic uncertainty in the model predictions, depending on which particular realization of the low-order multipoles is selected. Thus there is some variance in the *predicted* amplitudes of low-ℓ signals, even with a very specific model. Since our Universe represents only a single instance, the measured low-order multipoles could fall anywhere in the range projected by various realizations of the model. This range is cosmic variance. It is worth going on to make the obvious point that, in the range of ℓ dominated by cosmic variance, further improvements in the accuracy of observational results will not tighten constraints on the cosmological model or on cosmological parameters. We shall see that the results from the COBE satellite, launched nearly twenty years ago, have determined the low-order multipoles to adequate accuracy. Subsequent improvements in precision have resulted from better control of systematics and foregrounds.

$\ell \gtrsim$ 100. On scales below 2 degrees, the rich phenomena described so ably by Wayne Hu come into play. Density perturbations on the surface of last scattering produce a series of acoustic peaks in the CMB power spectrum, with the fundamental at $\ell \sim 200$. The

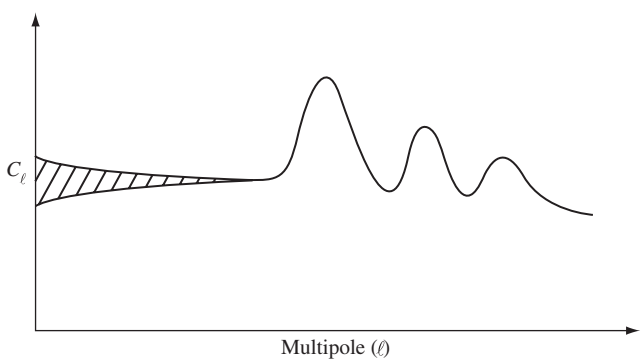

FIGURE 1.3. A cartoon version of the CMB power spectrum showing the various physical mechanisms discussed above.

physics explaining these peaks is quite well understood (but see some cautionary remarks by Wayne Hu), or at least well enough understood to give us reasonable confidence in the angular scale and relative amplitude of these acoustic peaks. These quantities in turn depend in predictable ways on crucial cosmological parameters such as the curvature, k, the baryon density, $\Omega_{\rm b}$, the Hubble parameter and so on. In this section and the next, we will explore some of the constraints on such parameters set by the observations of the acoustic peaks.

$\ell \gtrsim$ 500. As ℓ increases, the amplitude of the acoustic peaks is damped or cut off by two processes, Silk damping and the averaging of fluctuations of different sign resulting from the finite thickness of the last scattering surface. Both are described in Hu's article. Here, I use the latter effect to fix an approximate angular scale for the onset of this damping: the rough scale is given by $\theta \sim \frac{\Delta z}{z}(2^\circ)$, or something like 10 arc minutes, with $\Delta z \sim 150$ for the thickness of the last scattering surface. Thus these effects appear predominately at large values of ℓ.

In this section, I will treat measurements of the power spectrum on scales $\ell \lesssim 300$. That range includes the predicted plateau at low ℓ and the first acoustic peak, with emphasis on the angular scale of the latter. Although its amplitude owes primarily to local effects, rather than cosmological ones, I will also describe the dipole moment of the CMB. In Section 1.4, I will treat measurements of the entire series of acoustic peaks, evidence for Silk damping and the overall slope of the initial power spectrum in order to look for small departures from the $n_{\rm S} = 1$ expectation.

1.3.2 Some observational issues

I recognize that the division between $\ell \lesssim 300$ and $\ell \gtrsim 300$ is slightly artificial. It is inspired in part by the difference in observational techniques generally employed in these two regions of angular scale. Roughly speaking, as we shall see, the low ℓ measurements must be made from space or at least above a substantial fraction of the Earth's atmosphere, whereas ground-based measurements are possible at the smaller angular scales, $\ell \gtrsim 300$.

1.3.2.1 Comparative measurements

Unlike the case of spectral measurements discussed in Section 1.2 isotropy measurements can be made *comparatively*; absolute calibration is less pressing. The basic technique is to compare measurements made at two different positions in the sky. This can be done by fixing the apparatus and allowing the sky to pass overhead; by scanning the sky rapidly; or by beam switching, that is, comparing the difference in intensity observed by either two antennae pointed in different directions, or two beams from a single antenna employing a moving secondary or two feed horns. In practice, some or all of the techniques are combined to further reduce systematic errors. But the basic point is that anisotropy measurements can be made on the comparative basis, which is the main explanation for their far greater sensitivity than absolute measurements. MicroKelvin (μK) accuracy, rather than milliKelvin (mK) accuracy is required and can be attained. The Differential Microwave Radiometer (DMR) instruments aboard COBE, shown below in Fig. 1.4a, provide a good example. Matched pairs of microwave antennae, with their optical axes separated by 60°, fed signals into a comparative radiometer. Thus the difference in observed intensity from the sky was measured at the switching frequency, which was kept high enough to avoid gain drifts in the receiver. In addition, as the satellite rotated about its body axis, the two directions scanned circles in the

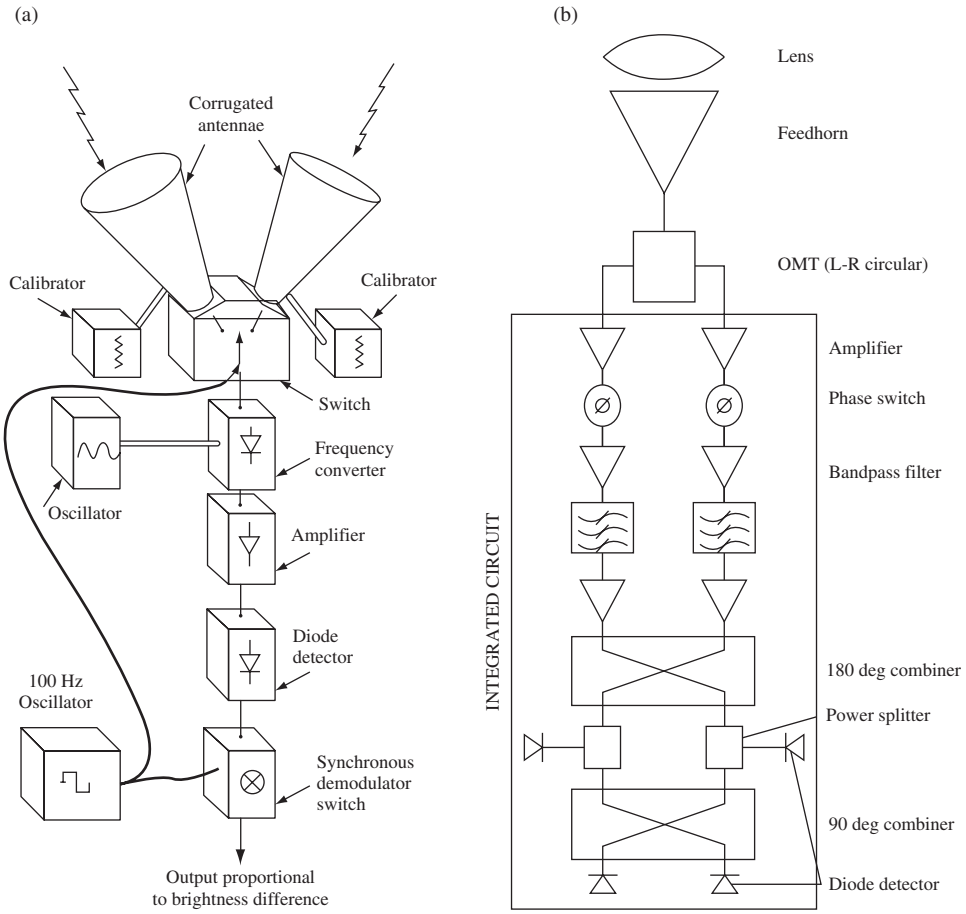

FIGURE 1.4. (**a**) Block diagram of the DMR instrument on COBE, showing symmetrical horns and calibrators. (From COBE slide set. NASA image.). (**b**) A pseudo-correlation receiver for polarization measurements (diagram from Todd Gaier). Signals from the horn are split in an orthomode transducer (OMT), amplified and recombined in such a way as to produce two linearly polarized outputs. Both polarizations pass through both amplifiers; thus variations in the gain average out.

sky. The annual passage of the satellite around the sun shifted the scanning directions, so that a complex, interlocking series of measurements of temperature differences in the CMB could be made. Unlike DMR, the Planck receivers will compare measurements of the sky with a controlled temperature cold load. The Planck satellite thus relies on the rapid rotation of the spacecraft to control gain drifts.

1.3.2.2 Calibration

To ensure that measured differences can be expressed in temperature terms, to compare to the 2.725 Kelvin background temperature, good calibration is required. It need not be at the same precision as the calibration required for absolute temperature measurements; they require precision of $\lesssim 1$ mK for calibration. In FIRAS this accuracy, or $\sim 4 \times 10^{-4}$ of the signal, was achieved. The same relative precision in calibration would allow us to reach 1–2 μK for a comparative measurement.

It is interesting to note that the most recent large-scale anisotropy experiments employ the COBE-measured dipole as the primary external calibration. The sensitivity of modern experiments, such as the Wilkinson Microwave Anisotropy Probe (WMAP) satellite to be discussed later, is so high that they can quickly detect the CMB dipole first accurately measured by COBE, and use it to calibrate the instruments. This turns out to be more accurate than the use of "point" sources, such as planets, as calibrators.

1.3.2.3 Other instrumental effects

To make accurate comparative measurements, we need to ensure that there are no instrumental offsets, and more importantly that instrumental offsets do not change with time. For instance, for observations made from beneath the atmosphere, the two directions being compared must lie at the same elevation, to avoid differential atmospheric emission. Equally, care must be taken to ensure that any signals diffracted into the antennae be independent of antenna alignment and preferably not time variable. Great care must also be taken to ensure that the gain and any loss in the receiver is symmetrical for the two directions being compared. In this regard, the use of a pseudo-correlation receiver, one that explicitly mixes the two signals, is much preferred. The design of such a receiving element is included as Figure 1.4b, and is explained further in the caption.

While eyes may glaze over with the discussion of some of these instrumental effects, controlling them is crucial to ensure the accuracy of CMB measurements. A measure of their importance can be seen in the series of long explanatory articles published as part of the analysis of the first three years of WMAP data (see below). Concern about instrumental effects also has an important impact on the *design* of CMB experiments, favoring those with a high degree of symmetry (as is true for WMAP).

1.3.3 Foreground emission

Even a perfect instrument faces the problem of foreground emission. The two major sources are the Earth's atmosphere and the Galaxy. I will defer discussion of a third source of foreground contamination, emission by radio galaxies and dusty galaxies, to Section 1.4.

1.3.3.1 Atmospheric fluctuations

If the atmosphere were an absolutely homogeneous, plane-parallel slab, as implicitly assumed in Section 1.2, atmospheric emission would simply cancel out if we compared two sky regions at the same elevation. But the atmosphere is turbulent, and turbulence introduces fluctuations, particularly on large angular scales. Indeed, it is the scale-dependence of atmospheric fluctuations that makes it extremely difficult to carry out CMB observations at $\ell \lesssim 300$ from the surface of the Earth.

First, let us look at the spatial dependence of atmospheric fluctuations. To first order, they can be modeled as a Kolmogorov process. For the conditions obtaining in the Earth's atmosphere, that predicts a power-law dependence of root-mean-square (rms) noise on linear scale as $d^{\frac{5}{6}}$. If we then make the assumption that these fluctuations are carried past an antenna beam at a uniform velocity, the rms fluctuations will scale with *time* to the same power. Some observational work by Smoot *et al.* (1987) finds time variation with an index of approximately 0.7, in rough agreement with the simple predictions. If we again assume a constant wind speed, we can back-calculate from this result to show that the amplitude of fluctuations scales with their length as $d^{0.7}$. Thus a high price is paid for observations on large angular scales. It is true that atmospheric fluctuations can