Chapter 2

Telescope, Receiver, and Radiometry

In this chapter, we discuss the telescope, optics, and receiver used to carry out the blazar monitoring program. We also describe the radiometry and calibration procedures employed to make the measurements.

For a monitoring program of this size and cadence, making efficient use of telescope time is critically important. Although we were extremely fortunate to have use of the Owens Valley Radio Observatory (OVRO) 40 m telescope 100% of the time from the start of the program until mid-2011, and six days a week after that, our cadence of 1500 sources every three days requires careful planning. Fortunately, the majority of our sources are reasonably bright—more than 50 mJy at 15 GHz—so the sensitivity requirements of the program are relatively modest. This has allowed us to optimize for rapid observations and easily repeatable measurements rather than scrabbling for sensitivity at all costs. Still, a full understanding of the behavior of the telescope and receiver and careful measurement and calibration are essential.

2.1 The Hardware

We begin by introducing the OVRO 40 m Telescope, its optics, and the Ku-band receiver used for this program.

2.1.1 The OVRO 40 m Telescope

The OVRO "40 m" telescope is actually a 130-foot-diameter f/0.4 parabolic reflector with approximately 1.1 mm rms surface accuracy on an altitude-azimuth mount. The telescope is located on the floor of the Owens Valley near Big Pine, California, at $37^\circ 13' 53\overset{''}{.}7$ N latitude, $118^\circ 16' 53\overset{''}{.}83$ W longitude, and 1236 m elevation (Pearson 1999). Construction of the 40 m telescope was completed in 1966. The telescope has been used with several different receivers since then, at frequencies as high as 45 GHz, where the surface accuracy of the dish becomes a serious limit on antenna efficiency.

2.1.1.1 Telescope Control Systems

The 40 m telescope is controlled by a computer control system that provides a user interface, executes schedules, controls the drive system servos, and records radiometer output and housekeeping data for later analysis. From well before the inception of this monitoring program until August 2010, these services were provided by a control system running a Digital Equipment Corporation VAX microcomputer with user interface functions on another VAX microcomputer connected via local-area network. This control system will be henceforth referred to as the *VAX control system*. Although the VAX control system and hardware had performed admirably since its installation in the early 1990s, increasingly frequent hardware failures and concern about replacement components and maintenance led to the design and implementation of a replacement system.

On 11 August 2010, the VAX control system was permanently disconnected and a new control system designed and written by Martin C. Shepherd, henceforth the *MCS control system*, was put into operation. The MCS control system runs on a personal computer using a real-time variant of the Linux operating system. In addition to operating on more easily replaced hardware, the new control system makes use of the vastly increased capabilities of modern computer hardware to log data at a greatly increased rate and to provide a more sophisticated scheduling system.

Although the control system was replaced, the receiver, digitizer, and drive hardware were unchanged. As a result, in large part the observing methods were unaffected by this change. The most significant impact will be discussed in section 3.1 where we describe the software tools and reduction scripts, which were rewritten to work with the new data format.

2.1.1.2 Mount and Drive System

The 40 m telescope is mounted on an altitude-azimuth mount. Azimuth is measured from North through East with 0° at due North. The telescope can slew through 425°, with an overlap region in the northwest quadrant between -90° and $+335^\circ$. In elevation, the telescope can be pointed from 11.5° above the horizon to about 10° past zenith. In practice, however, the control system limits the maximum elevation to 90°, and observations are normally only made between 30° and 70° elevation to avoid excessive airmasses at low elevations and because the drive system has difficulty matching the sidereal rate near zenith. The telescope can be slewed at a maximum rate of about 15° per minute in both azimuth and elevation, but can only track a moving source at half this rate or less.

In a small azimuth range pointed due South (azimuth 180°), the telescope can be tilted down to 7° elevation. In this "service position," a ladder in one of the focus support legs enables access to the prime focus. This position is used for service, maintenance, and calibration procedures that require access to the receiver or optics. In particular, hot/cold load Y-factor measurements are performed in this position (see section 2.2.1.1).

Table 2.1. List of thermometers instrumenting the 40 m telescope

Label	Typ. Val.	Units	Description
T_{RX}	24	°C	Ambient section of receiver enclosure
T_{HEMT}	12	K	HEMT LNA temperature
$T_{70\ \mathrm{K}}$	66	K	70 K stage
$T_{15\ \mathrm{K}}$	11	K	15 K stage
T_{plate}	21	°C	Cold plate temperature
T_{switch}	80	K	Dicke switch temperature
T_{hot}	300	K	Hot load temperature during Y-factor measurement, approximate outdoor temperature otherwise
T_{backend}	26	°C	Receiver backend

2.1.1.3 Tilt and Temperature Monitoring

The 40 m telescope is equipped with two orthogonal tilt meters located in the teepee of the telescope in the alidade above the azimuth bearing. These are carefully aligned with the telescope axes and are referred to as the aft/forward and left/right meters. The purpose of these sensors is to monitor for tilt of the drive system relative to the topocentric coordinate frame due to gravity or wind. These tilts are in the range of up to a few arcminutes at most. The tilt sensor readings are used in the pointing model (see section 2.2.2.1) with a scale parameter to compensate for errors in their output calibration.

A set of thermometers at the prime focus monitor temperatures related to the receiver and its support electronics. Table 2.1 lists the thermometers and their purposes. These are primarily used to monitor for problems in the receiver, except for T_{hot}, which is used during the hot/cold load Y-factor measurement (see section 2.2.1.1).

2.1.1.4 Weather Measurements

The weather is an important consideration for our observations, most critically because in moderate winds the telescope cannot be accurately pointed, and high winds can even threaten the safety of the telescope. From the beginning of the monitoring program until the transition to the MCS control system, a simple digital weather station equipped with an anemometer reported the instantaneous wind speed to the control system. In 2009, a Columbia Weather Systems[1] Capricorn 2000EX weather station was installed. The data from the new weather station were logged for offline use when using the VAX control system. With the switch to the MCS control system, which uses the Capricorn 2000EX for real-time weather monitoring, the old weather station was retired. The Capricorn 2000EX performs instantaneous and peak gust wind measurements, as well as ambient temperature, precipitation, barometric pressure, and relative humidity measurements.

In high wind situations, wind loading on the telescope could exceed the power of the drive system, potentially leading to damage to the telescope. To prevent this, in high wind conditions, the observing program is suspended and the telescope is steered to the "stow position" at about 180° azimuth, 90° elevation. In this

---

[1] `http://www.columbiaweather.com`

position, the cross section of the telescope to the wind is minimized, so it is the safest orientation in high winds. With the original weather station, such a "wind stow" was triggered when the instantaneous wind speed exceeded 25 mph (11.2 $\mathrm{m\,s^{-1}}$). Using the Capricorn 2000EX, a wind stow is triggered by a 60 min peak gust speed of 25 mph (11.2 $\mathrm{m\,s^{-1}}$) or an instantaneous wind speed of 20 mph (8.9 $\mathrm{m\,s^{-1}}$. The wind stow is maintained for at least 60 min and until the instantaneous speed has fallen below 18 mph (8.0 $\mathrm{m\,s^{-1}}$) and the 60 min peak gust has fallen below 22 mph (9.8 $\mathrm{m\,s^{-1}}$).

In addition to triggering wind stows, the wind speed is also used to identify periods when the pointing of the telescope was degraded due to high winds. This is discussed in section 3.2.1.2.

2.1.2 Optics

At the prime focus of the 40 m telescope, two symmetric off-axis corrugated horn feeds are installed inside the receiver cryostat. Coupled to the parabolic reflector of the 40 m telescope, this produces a pair of approximately Gaussian beams with $157''$ FWHM, separated by $12\overset{\prime}{.}95$ on the sky. We refer to these two beams, somewhat arbitrarily, as the "antenna" beam and the "reference" beam, or *ant* and *ref*. The beam separation is in the azimuthal direction, and because the beams are offset symmetrically from the optical axis, they are always located at equal elevations.

Figure 2.1 shows a schematic view of the optics and the waveguide section of the receiver before the low-noise amplifier. In this section, both the *ant* and *ref* signal paths are identical. After the feed, a dielectric waveguide polarizer selects left-hand circular polarization (LCP); because the received radiation is reflected from the telescope, this corresponds to right-hand circular polarization (RCP) on the sky (M. W. Hodges, personal communication; Moffet 1973).[2] As a result, linearly polarized sources of all orientations may be monitored in total intensity. The signal in each then passes through a circular-to-square waveguide transition, through a 30 dB directional coupler, and then into the Dicke switch. The Dicke switch common port and the directional couplers' ports each pass through a transition to a coaxial connector that connects to the rest of the receiver.

On the *ant* side, the directional coupler connects to the calibration noise diodes with a 30 dB reduction of the diode signal. The *ref* directional coupler is simply terminated and is included only to maintain symmetry between the two signal paths. The Dicke switch port connects to the HEMT low-noise amplifier. The calibration diodes and the receiver are discussed in section 2.1.3.

---

[2] Here, we adopt the Institute of Electrical and Electronics Engineers (IEEE) circular polarization convention that RCP corresponds to a clockwise temporal rotation (at a fixed position) of the electric vector from the point of view of the source, i.e., when looking in the direction of propagation.

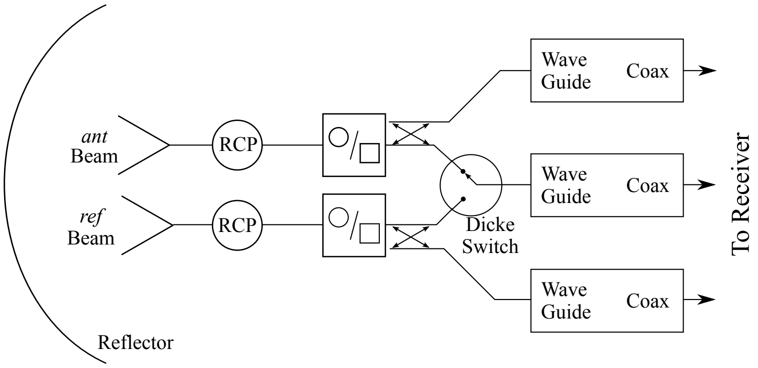

Figure 2.1. Optics and waveguide section block diagram.

2.1.2.1 Aperture Efficiency

The power received by a radio telescope with effective receiving area A_e sensitive to a single polarization is

$$P = \frac{1}{2}\, A_e\, S_\nu\, \Delta\nu, \tag{2.1}$$

where S_ν is the incident flux density and $\Delta\nu$ is the receiver bandwidth. If the antenna has a physical aperture area A_p, the aperture efficiency, η_a, is defined by

$$A_e = \eta_a\, A_p. \tag{2.2}$$

The aperture efficiency can be factored into contributions from a number of different causes. For example, if we combine the effects of feed illumination, spillover, and blockage into η_i, and quantify the effect of phase errors due to surface irregularities with η_p,

$$\eta_a = \eta_i\, \eta_p. \tag{2.3}$$

For an unresolved point source—a source of angular extent much smaller than the beam of the telescope—the specific intensity is effectively a delta function in angle. The response to a point source of flux density S_ν will then be

$$P = \frac{1}{2} A_e\, \Delta\nu \iint S_\nu\, \delta(\theta - \theta_0)\, \delta(\phi - \phi_0)\, B(\theta, \phi)\, d\Omega = \frac{1}{2} A_e\, \Delta\nu\, S_\nu\, B(\theta_0, \phi_0), \tag{2.4}$$

where $B(\theta, \phi)$ is the normalized antenna gain and (θ_0, ϕ_0) is the position of the source in the beam. Assuming this is centered, $B(0, 0) = 1$, so $P = A_e\, \Delta\nu\, S_\nu$. The antenna temperature, T_a, due to this point source is the temperature of a blackbody filling the aperture that gives the same response. The response to the blackbody is

$$P = \frac{1}{2} A_e\, \Delta\nu \iint I_\nu^{RJ}\, B(\theta, \phi)\, d\Omega = \frac{1}{2} A_e\, \Delta\nu\, I_\nu^{RJ}\, \Omega_a, \tag{2.5}$$

Table 2.2. Aperture efficiency measurement results

Date	η_a	Source(s)
2008-11-08	0.240	DR 21, NGC 7027, 3C 286
2009-03-11	0.255	3C 286
2009-05-19	0.247	3C 48
2009-07-07	0.261	3C 286
2009-08-10	0.275	3C 286
2009-09-10	0.264	3C 286
2009-11-11	0.269	DR 21, NGC 7027, 3C 48
2010-02-08	0.247	DR 21, NGC 7027, 3C 48
2010-04-26	0.260	3C 48
Mean	0.258 ± 0.004	

where $\Omega_a = \iint B(\theta, \phi)\, d\Omega$ is the beam solid angle. In the Rayleigh-Jeans limit ($h\,\nu \ll k_B\,T_a$), the specific intensity is proportional to temperature:

$$I_\nu^{RJ} = \frac{2\,k_B\,T_a}{\lambda^2}, \tag{2.6}$$

where k_B is Boltzmann's constant and λ is the wavelength. By equating the two detected powers from equations (2.4) and (2.5), we find

$$2\,k_B\,T_a = \frac{S_\nu\,\lambda^2}{\Omega_a}. \tag{2.7}$$

By the antenna theorem (e.g., Rohlfs & Wilson 2000),

$$A_e\,\Omega_a = \lambda^2, \tag{2.8}$$

so

$$2\,k_B\,T_a = S_\nu A_e = S_\nu\,\eta_a\,A_p. \tag{2.9}$$

We use equation (2.9) to measure the aperture efficiency of the 40 m telescope. Solving for η_a, we have

$$\eta_a = \frac{2\,k_B\,T_a}{A_p\,S_\nu} \tag{2.10}$$

where T_a is the measured antenna temperature for an unresolved point source of known flux density S_ν. In practice, we determine T_a by comparing the detected signal to a measurement of the CAL diode, whose equivalent noise temperature we know from the Y-factor tests described in section 2.2.1.1. This both converts the digitizer units (DU) to K and corrects for nonlinearity because the CAL diode measurement is affected by nearly the same amount of gain compression as the source measurement.

In table 2.2, we tabulate the measurements of η_a made during this program. Combining these and estimating uncertainty from the sample standard deviation, we find the aperture efficiency $\eta_a = (0.258 \pm 0.004)$. This relatively low aperture efficiency is due to deliberate underillumination of the dish by the feed—for monitoring observations of a large sample of objects aiming at flux density measurements repeatable to within

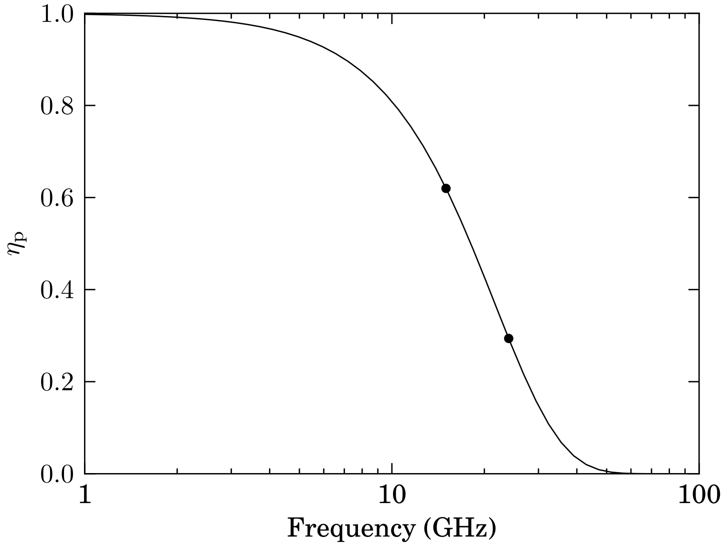

Figure 2.2. Predicted efficiency factor η_p calculated from equation (2.11). Points indicate 15 GHz (0.62) and 24 GHz (0.29) values.

a few percent we must consider the trade-off between aperture efficiency and pointing accuracy. Underillumination of the antenna increases the beamwidth and reduces susceptibility to pointing errors relative to more fully illuminating the antenna, in addition to reducing exposure to thermal noise from ground spillover. Experience has shown that we are operating at close to the optimum illumination for the most efficient use of the telescope at 15 GHz: increasing the aperture efficiency gains little because the thermal noise is already acceptably low for observing the objects in our monitoring sample.

2.1.2.2 Surface Accuracy

The surface of the 40 m telescope is composed of 852 individually adjustable panels, each with an surface accuracy of about 0.36 mm. After adjustment to match a parabolic figure at about 50° elevation, the total surface accuracy is about 1.1 mm rms. The *Ruze formula* predicts the reduction of the aperture efficiency at frequency ν due to surface errors with $\mathrm{rms}\ \epsilon$ to be

$$\eta_p = e^{-(4\pi\,\epsilon\,\nu/c)^2}, \tag{2.11}$$

where c is the speed of light. Figure 2.2 shows the predicted values for η_p at various frequencies. At 15 GHz, $\eta_p = 0.62$. This accounts for a significant factor in the total aperture efficiency. To obtain the observed aperture efficiency $\eta_a = 0.258$, the illumination and blockage factors must amount to $\eta_i = \eta_a/\eta_p \approx 0.42$.

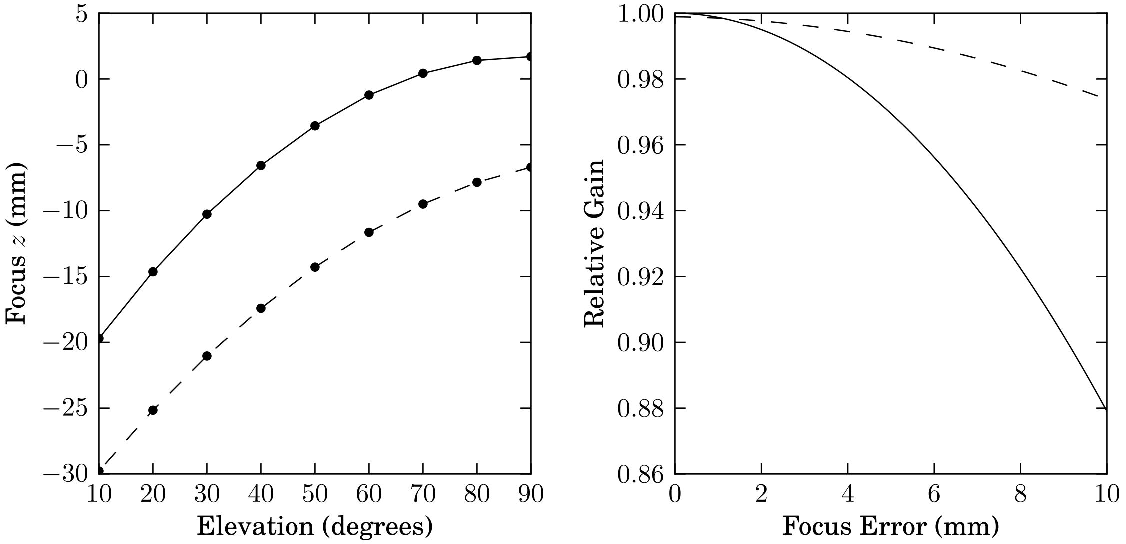

Figure 2.3. *Left:* Focus curve used during observations to predict optimum focus as a function of elevation. Points indicate the look-up table values, which are linearly interpolated as shown. The solid (dashed) line indicates the curve before (after) the shift in April 2010. *Right:* Relative gain that results from a focus error before (solid) and after (dashed) the shift. Measured flux densities are divided by the value of this curve to compensate for use of the simple focus curve rather than the more accurate model.

2.1.2.3 Antenna Gain and Focus

When the 40 m telescope moves in elevation, gravity deforms its surface, changing the antenna gain and focus location. The entire feed/receiver system can be moved along the optical axis to adjust the focus. The optimum focus position as a function of elevation is measured about once per year, but has not been found to vary significantly except when the receiver has been removed and reinstalled during maintenance. Due to thermal effects, the optimum focus also varies slightly between day and night operation and with the angle between the telescope structure and the Sun.

In normal operation, the focus is set before each observation procedure using a polynomial fit to the measured optimum focus as a function of elevation using a linearly interpolated look-up table. An example focus curve is shown in the left-hand panel of figure 2.3 with the plotted values given in table 2.3. The focus models do not change frequently. From 2008 until April 2010, measurements indicated there was no need to modify the focus curves. In April 2010, the receiver was removed from the telescope for maintenance and when reinstalled, an offset of nearly 1 cm was found in the optimum focus positions was observed, so a new focus model was determined.

During data calibration, a more complicated focus model that includes solar elongation and elevation is evaluated and a correction is applied to account for the focus error relative to that model. The ideal focus

Table 2.3. Focus curve values plotted in the left-hand panel of figure 2.3

Elevation (°)	z (Before) (mm)	z (After) (mm)
10	-19.70	-29.78
20	-14.64	-25.16
30	-10.27	-21.04
40	-6.57	-17.42
50	-3.56	-14.29
60	-1.22	-11.65
70	0.43	-9.50
80	1.41	-7.85
90	1.70	-6.70

Table 2.4. Polynomial coefficients for the focus models (before and after April 2010)

	Before			After		
n	a_n	b_n	c_n	a_n	b_n	c_n
0	7.02	—	—	-1.99	—	—
1	0.0102	-0.0677	-0.0355	-0.0241	-0.061	-0.0521
2	-0.00366	0.000209	0.000189	-0.00327	0.000156	0.000244

position, z, is given by

$$z = \sum_{n=0}^{N_a} a_n(90^\circ - \theta)^n + \sum_{n=1}^{N_b} b_n(90^\circ - \theta_\odot)^n + \sum_{n=1}^{N_c} c_n{\zeta_\odot}^n, \tag{2.12}$$

where θ is the source elevation, $\theta_\odot$ is the solar elevation, and $\zeta_\odot$ is the angular distance on the sky between the source and the sun (all measured in degrees) and a_n, b_n, and c_n are polynomial coefficients. These values are tabulated in table 2.4. The correction is calculated from a polynomial "focus miss" model, as shown in the right-hand panel in figure 2.3 with coefficients given in table 2.5. The focus miss model also changed significantly in April 2010, with the relative gain becoming much less sensitive to focus errors. This focus model and the parameters were developed and measured by Walter Max-Moerbeck by measuring the optimum focus position for point sources at many elevations and times of day.

Even with the focus adjustment and correction, the gain of the telescope is found to vary with elevation. This is due to reduced antenna gain as the reflector deforms under its own weight as it slews relative to the vertical. The surface of the reflector was set to provide an optimum parabolic surface at about 50° elevation (Pearson 1999). Figure 2.4 shows a 5th-order polynomial fit to the relative gain as a function of elevation, measured by tracking 3C 286 as it moved from about 20° to 80° elevation. The polynomial

Table 2.5. Polynomial coefficients of the focus miss curve (before and after April 2010)

n	c_n (before)	c_n (after)
0	1.0	0.999
1	-9.84×10^{-5}	-1.81×10^{-4}
2	-1.20×10^{-3}	-2.33×10^{-4}

Table 2.6. Gain curve polynomial coefficients (before and after April 2010)

n	c_n (before)		c_n (after)	
0	7.158	$\times 10^{-1}$	8.464	$\times 10^{-1}$
1	4.122	$\times 10^{-3}$	−1.210	$\times 10^{-2}$
2	7.457	$\times 10^{-4}$	1.744	$\times 10^{-3}$
3	−2.913	$\times 10^{-5}$	−6.030	$\times 10^{-5}$
4	3.752	$\times 10^{-7}$	8.270	$\times 10^{-7}$
5	−1.689	$\times 10^{-9}$	−4.106	$\times 10^{-9}$

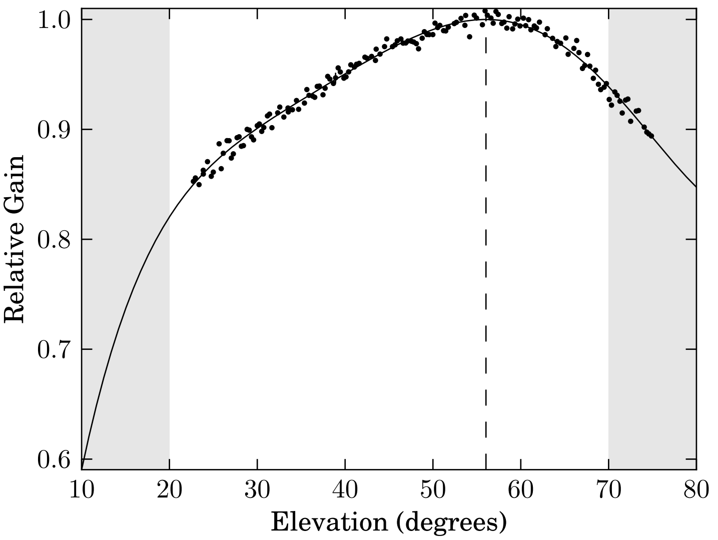

Figure 2.4. Example antenna gain curve plotting the relative peak gain as a function of elevation. Curve is a 5th-order polynomial fit to data collected from 3C 286 observations at a range of elevations on 09 March 2011. The coefficients of this polynomial are given in table 2.6 in the “After” column. Grey areas indicate regions where observations are not normally permitted and where the gain curve fit may be unreliable. Points indicate the data used for the fit, including measured uncertainties. The dashed line indicates the peak at 56.0° elevation.

coefficients for the gain curves used before and after the receiver was taken down and reinstalled in April 2010 are shown in table 2.6. Flux density measurements are corrected for this gain variation by dividing the observed flux density of a source by the value of the gain curve at the elevation of the observation. It is important to note that this correction is based on measurements of the response at the peak of the beam and is only appropriate for point-source observations like those that make up this observing program.

2.1.2.4 Beam Map

The telescope beam is characterized by the normalized power pattern $B(\theta, \phi)$, which gives the power response of the telescope to a uniform plane wave incident from direction (θ, ϕ) relative to the peak response $\max\{B\} = 1$. The two feeds at the 40 m prime focus project two symmetric beams on the sky, *ant* and *ref*,

Table 2.7. Properties of a few point sources suitable for beam mapping

Name	RA (J2000) ($^{\rm h\ m\ s}$)	Dec (J2000) ($^{\circ}\ '\ ''$)	Size[a] ($''$)	Flux Density (Jy)
3C 48	01 37 41.3	+33 09 35	<1	1.72
3C 147	05 42 36.1	+49 51 07	<1	2.65
3C 286	13 31 08.3	+30 30 33	<5	3.44
3C 295	14 11 20.5	+52 12 10	4	1.61

Source: Flux densities and angular sizes are from Baars et al. (1977).

[a] Angular sizes specified at 1.4 GHz.

Table 2.8. Results of fitting Gaussian components to the beam center scan in figure 2.5

Beam	Amplitude (DU)	ϕ ($'$)
ant	76.1	−0.24
ref	−82.2	12.73

with an angular separation Ψ. We decompose the beam response into separate terms for each beam, that is

$$B(\theta, \phi) = B_{\rm ant}(\theta, \phi - \frac{\Psi}{2}) + B_{\rm ref}(\theta, \phi + \frac{\Psi}{2}). \tag{2.13}$$

Because of the identical construction and symmetric placement of the feeds, the individual beams are expected to have very similar responses relative to their centers, with any deviations likely to be mirror-symmetric between the beams. Unless otherwise specified, when we describe properties "the beam," we refer to one of the two offset beams, $B_{\rm ant}$ or $B_{\rm ref}$.

The coordinates (θ, ϕ) are given relative to the optical axis of the telescope. Neglecting misalignment with the mount, the θ axis is the same as the elevation axis and the ϕ axis measures angle along the great circle on the sky that is tangent to the azimuth axis. These coordinates are properly measured in the spherical geometry of the sky. However, for the very small angular extent ($\lesssim 15'$) of the 40 m beam, we can safely treat the coordinates as Cartesian, remembering that the scale factor between ϕ and the mount's azimuth coordinate varies with elevation.

To measure the beam response, we use an unresolved astronomical source to sample the response of the telescope at various angular offsets. For the 40 m with a beam FWHM $\approx 2\overset{\prime}{.}6$, a source must have an angular size $\ll 1'$ to be unresolved and regarded as a point source. Table 2.7 lists the properties of a few suitable sources.

In figure 2.5, we show the result of 50 min continuous scanning in ϕ across a source (3C 295) positioned at the elevation center of the beams. Each scan spanned $\pm 2^{\circ}$ around the source, but only the center region is plotted. After removing the median background, we fitted an independent Gaussian profile with a fixed $157''$ FWHM and free amplitude and ϕ position to each beam. The fitting results are shown in table 2.8. The separation between the beams was thus measured to be $12\overset{\prime}{.}97$. This is sufficiently close to the previous value of $\Psi = 12\overset{\prime}{.}95$, reported by Bustos (2008), that we continue to quote the old value for continuity.

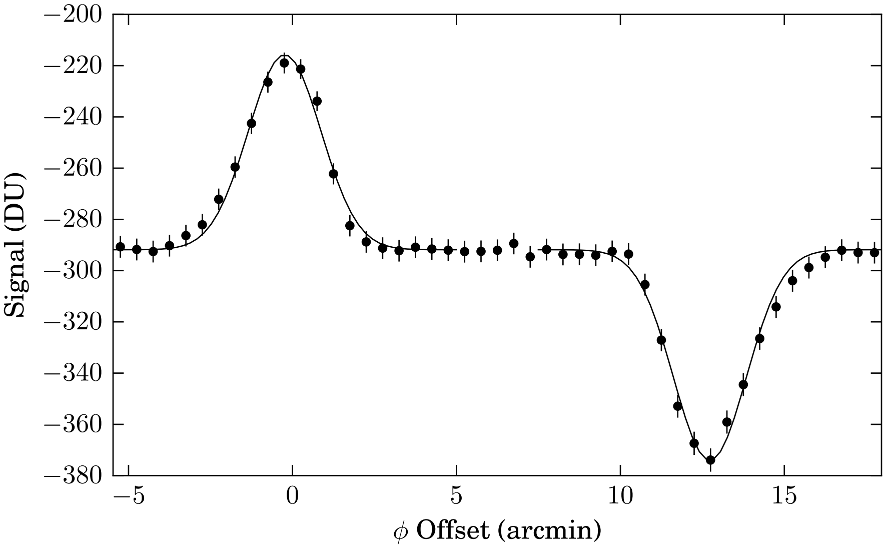

Figure 2.5. Binned switched signal (*ant* − *ref*) from 50 min of continuous azimuth scans through the elevation center of both *ant* and *ref* beams, measured on 31 Aug, 2011, using 3C 295. Lines are individual fixed-beamwidth Gaussian fits to the data spanned by the plotted line. The separation between the two beams is 12.′97 on the sky.

In July 2011, we measured the beam response using 3C 286, 3C 295, and 3C 48. FLUX procedures were performed on a triangular grid with 45″ spacing between centers. Because the beam changes as the telescope changes elevation, we restrict the measurements to elevations near 45°. Scans at constant elevation offsets from the source position were performed as sources rose or set through approximately 40°–50° elevation. To cover the region with radius $4'$ requires 109 pointings, or about 2 hours of observing, which we split into alternate elevation rows due to the elevation constraint. Several repetitions of this procedure are averaged to reduce noise, then contours are computed and plotted using the Matplotlib[3] `tricontour` routine to linearly interpolate between data points on the triangular grid.

Because FLUX procedures sample the source in both the *ant* and *ref* beams, this procedure measures the average of the two beams. This resulting average beam is shown in figure 2.6. To separate the beam responses, we also perform the same scans with an additional ϕ offset equal to the beam separation, Ψ.

Using this offset, a FLUX procedure now measures blank sky during the A and D segments and the source through the *ref* beam in the B and C segments. This, of course, reduces the on-source integration time by a factor of two, so results in a reduction in the gain of the FLUX procedure. Using an offset of $-\Psi$ similarly allows measurement of the *ant* beam alone. The separated beams, measured using 3C 48 are shown in figure 2.7. It appears that a slight inclination relative to the intended azimuthal separation is present. Elliptical Gaussian beams were fitted to the *ant* and *ref* beams individually and are plotted using the same (θ, ϕ) grid and contours in figure 2.8. Residuals from the fits are shown in figure 2.9 and the fit

---

[3] `http://matplotlib.scipy.org`

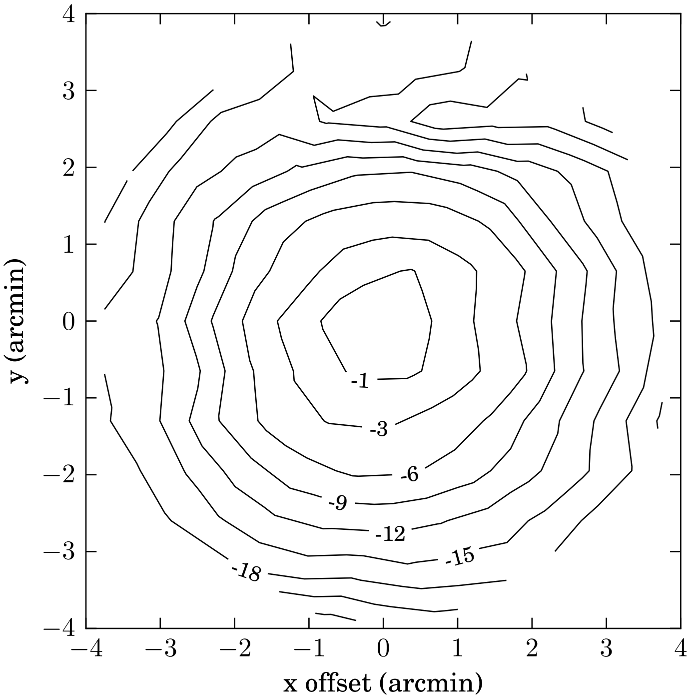

Figure 2.6. Normalized beam response for the average of the *ant* and *ref* beams, measured using 3C 286 and 3C 295. Contours are in dB relative to the peak. This depicts the effective beam for a FLUX procedure (neglecting the negative reference field lobes that are ±12$'.95$ away).

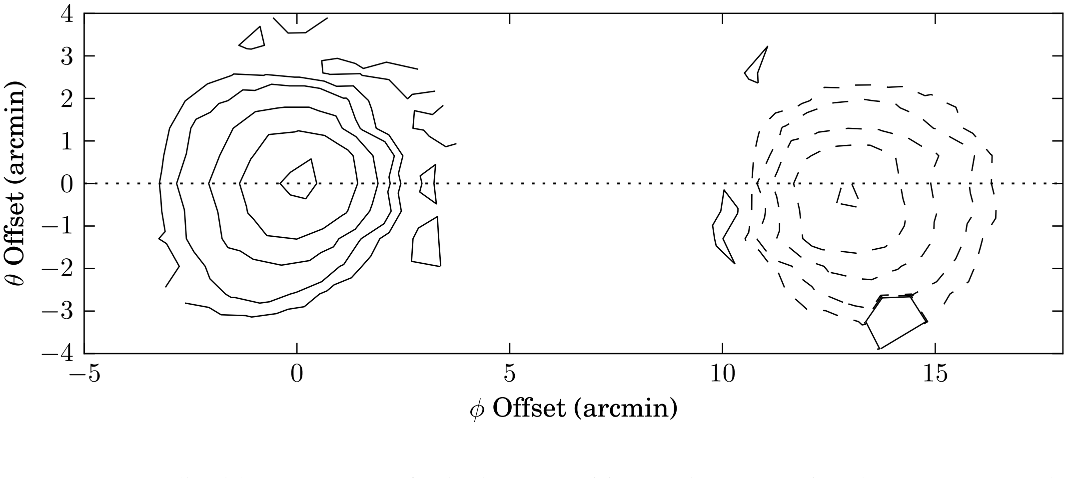

Figure 2.7. Normalized beam response for both *ant* (positive) and *ref* (negative) beams. Measured using 3C 48. Contours are 0, ±5%, ±10%, ±25%, ±50%, and ±90% of the *ant* peak, with dashed contours indicating negative values. The dotted line indicates pure azimuthal offset along which the beam separation is measured (see figure 2.5). Some inclination relative to this axis is apparent.

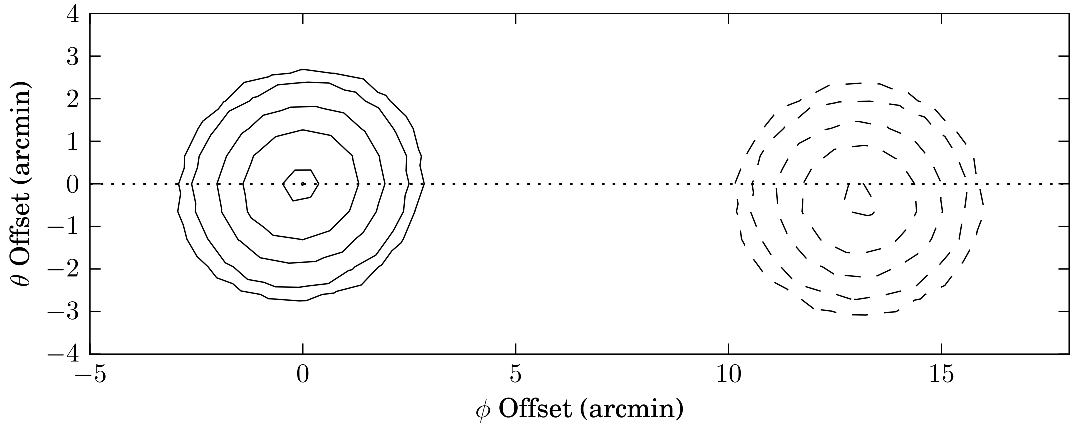

Figure 2.8. Elliptical Gaussian beam fits to the *ant* (positive) and *ref* (negative) beams. Contours are 0, ±5%, ±10%, ±25%, ±50%, and ±90% of the measured *ant* peak, matching the contours in figure 2.7.

Table 2.9. Elliptical Gaussian beam fit parameters

Parameter	*ant*	*ref*
Amplitude	1	0.99
Major FWHM	3′.07	2′.88
Minor FWHM	2′.57	2′.69
ϕ center	0	13′.10
θ center	0	−0′.28
Major axis inclination	−57°	−60°

Note: The Gaussian parameterization is that described in Leitch (1998).

parameters are shown in table 2.8. The results suggest that our nominal adopted beam FWHM of 157″ (2′.62) is underestimated, but this does not affect the observations in this program.

2.1.3 Receiver

A block diagram of the receiver is shown in figure 2.10. The receiver operates in the Ku band with a center frequency of 15.0 GHz. The receiver noise temperature is about 30 K, and the typical system noise temperature including receiver, cosmic microwave background (CMB), atmospheric, and ground contributions is about 55 K.

The receiver front end consists of a cooled (T≈80 K), low-loss ferrite RF Dicke switch followed by a cryogenic (T≈13 K) HEMT low-noise amplifier. This is followed by additional room-temperature amplifiers, a 13.5–16.5 GHz band definition filter, and an electronically controlled attenuator used to adjust the overall gain of the receiver. The signal is detected directly using a square law detector diode. The detected signal is digitized with a 16-bit analog-to-digital converter and then recorded.

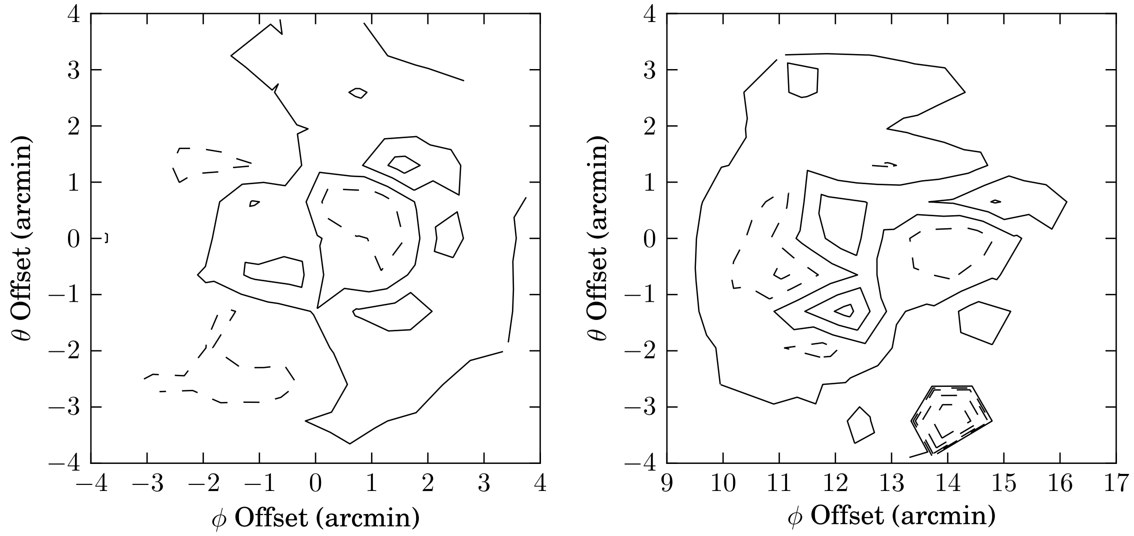

Figure 2.9. Residuals from fits to the beam map data shown in figure 2.7 for the *ant* (left) and *ref* (right). Contours show 0, ±5%, ±10%, and ±25% of the *ant* peak value.

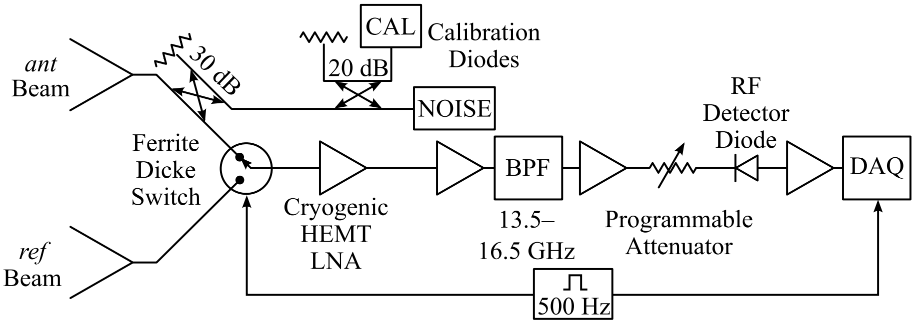

Figure 2.10. Block diagram of the Ku-band receiver.

2.1.3.1 Dicke Switching

In order to make the most efficient use of the telescope, a Dicke-switched dual-beam system is used (e.g., Rohlfs & Wilson 2000). The ferrite RF Dicke switch is switched at 500 Hz, alternately delivers the *ant* and *ref* beams to the receiver input. The radiometer output is integrated by an analog integrator circuit in each 1 ms half-period and then sampled. In software, the *ant* and *ref* samples are normally subtracted to produce the switched power, $\xi = P_{ant} - P_{refbeam}$. In some applications, such as when computing a nonlinearity correction, the average (or total) power $P = 0.5(P_{ant} + P_{ref})$ is required.

The most important benefit of Dicke switching is the removal of the large, slowly varying total power signal, which is made up of contributions from ground, atmosphere, and receiver thermal noise. Variations in the gain of the low noise amplifier cause variations in the large total power signal, and in addition the signals themselves vary slowly with time and with the position of the telescope. The resulting large variations in power limit the sensitivity of the receiving system, as discussed in section 2.1.3.3. Ground spillover, like gain variations, contributes directly to the system noise, but the effect is difficult to quantify due to the complexity of the far sidelobes of the telescope beam. Dicke switching removes or reduces these large slowly varying signals.

A second benefit of Dicke switching is the reduction of noise due to the rapidly varying atmosphere above the telescope. With a beam separation of $12\rlap{.}'95$, and for a water vapor scale height of 1.5 km, 75% of the total mass of water vapor seen by the telescope lies in the overlapping portions of the two beams. This fraction does not change substantially with scale height, dropping only to 72% (69%) for a water vapor scale height of 2 km (2.5 km). So Dicke switching reduces the effects of the varying atmosphere by about a factor of four.

A third benefit of Dicke switching is that the on-off measurement of the source against the reference allows the flux density of the source to be measured in a single pointing. This is much faster than the alternative strategy of scanning a single beam across the source. Additionally, because the source is near the peak of the beam response for the entire integration, the effective sensitivity is greater for the same integration time. This is at the cost of more stringent pointing requirements, since a mispointing will reduce the apparent brightness of the source. More details of the flux density measurement procedure are provided in section 2.2.2.2 below.

2.1.3.2 Bandwidth

The output of the receiver in response to a narrowband input signal varies depending on the frequency of that input. In general, this is a complicated function with peaks and valleys. However, the response is normally approximately zero except in some range of frequencies around the nominal frequency of the receiver. The width of this range is characterized by the *bandwidth* of the receiver. Qualitatively, the meaning of bandwidth is clear. However, there are several quantitative definitions, each useful for different calculations. Three common definitions are the half-power bandwidth, $\Delta\nu$, the noise bandwidth, $\Delta\nu_{\rm noise}$, and the radiometer

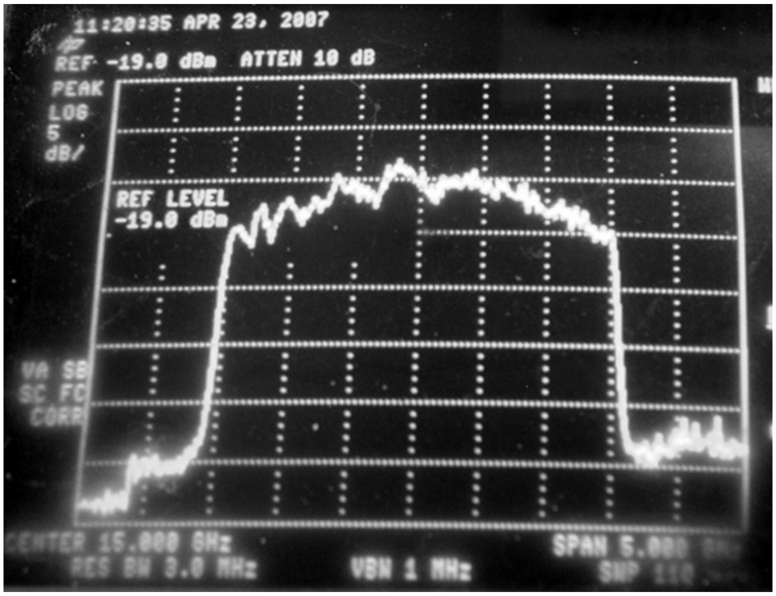

Figure 2.11. Photograph of spectrum analyzer sweep of the receiver response. Center frequency is 15 GHz, frequency span is 5 GHz (i.e., 500 MHz per division). Vertical scale is 5 dB per division. Reproduced from Bustos (2008).

reception bandwidth, $\Delta\nu_{\rm rec}$. For reasonably flat frequency responses, these bandwidths are of similar magnitude. In the literature, the nomenclature for these definitions varies, so care must be taken to determine what definition a particular author is using. Here, we adopt the convention used by Evans & McLeish (1977).

To compute these bandwidths, we begin with a spectrum analyzer trace of the receiver response. Figure 2.11 shows this response on a semilogarithmic plot, reproduced here from Bustos (2008). To work quantitatively, a piecewise linear approximation to the curve was estimated in the pass band between 13.5 and 16.5 GHz. This estimate is shown in figure 2.12 and the estimated values are tabulated in table 2.10.

Half-power bandwidth. The half-power, or 3 dB, bandwidth is the difference between the frequencies at which the receiver's power response is half that of the peak. If the ripple in the response is greater than 3 dB, the lowest and highest 3 dB points are used. This is the simplest bandwidth to measure and is frequently implied when a specific bandwidth definition is not given. Using the approximate bandpass data plotted in figure 2.12, the 40 m Ku-band receiver has a half-power bandwidth $\Delta\nu = 1.5$ GHz between 14.3 and 15.8 GHz.

Noise bandwidth. The noise bandwidth is the bandwidth of a hypothetical receiver with perfectly flat response, the same peak gain, and the same response to a wideband white noise input as the receiver in question. That is,

$$\Delta\nu_{\rm noise} \equiv \int_0^\infty \frac{G(\nu)}{G_{\rm max}} d\nu. \tag{2.14}$$

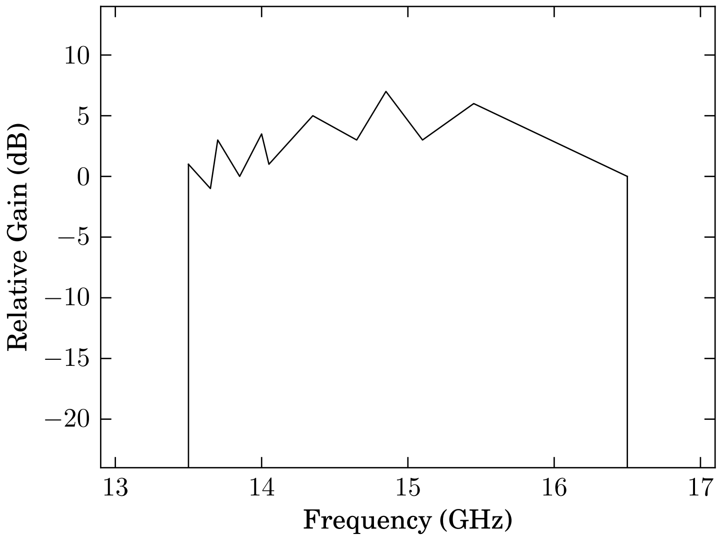

Figure 2.12. Piecewise linear approximation to the spectrum analyzer response shown in figure 2.11. Numerical values are listed in table 2.10. Attenuation is assumed infinite outside the 13.5–16.5 GHz band.

Table 2.10. Segment endpoints for the piecewise linear approximation to the measured receiver gain

ν (GHz)	G_{dB} (dB)
13.50	1.0
13.65	−1.0
13.70	3.0
13.85	0.0
14.00	3.5
14.05	1.0
14.35	5.0
14.65	3.0
14.85	7.0
15.10	3.0
15.45	6.0
16.50	0.0

Note: The approximation linearly connects these points in the semilogarithmic gain-frequency plane.

The noise bandwidth of the postdetection filter is a particularly important application of this definition. As discussed in section 2.1.3.3, it quantifies the postdetection circuit's contribution to the radiometer sensitivity. In this section, however, we compute the noise bandwidth of the receiver.

To compute the noise bandwidth from our approximate response, we must first convert the log-linear piecewise approximation into a linear-linear model. Normalizing for unity peak gain, the result is shown in the solid lines in figure 2.13. Note that our piecewise-linear model was expressed in dB, so it becomes a piecewise-exponential model rather than the piecewise-linear model shown in dashed lines. The distinction between the two is small, but we would slightly overestimate bandwidths by using the linear interpolation.

Although we could integrate numerically, it is straightforward to evaluate the integral analytically for one exponential segment of our model. For later use, we will evaluate the integral of an arbitrary nonzero power of the gain, p. If the kth segment connects (ν_0, G_0) with (ν_1, G_1), we have

$$I_k^{(p)} = \int_{\nu_0}^{\nu_1} \left(G_0 \, 10^{m(\nu - \nu_0)} \right)^p d\nu, \tag{2.15}$$

where m is the semilogarithmic slope of the segment. This has the solution

$$I_k^{(p)} = \frac{(\nu_1 - \nu_0)\,(G_1^p - G_0^p)}{p\ (\ln G_1 - \ln G_0)}. \tag{2.16}$$

Applying this to the data from table 2.10 gives $\Delta\nu_{\mathrm{noise}} = \sum_k I_k^{(1)}/G_{\mathrm{max}} = 1.37$ GHz.

As a test of the sensitivity of the result to errors in reading data points from figure 2.11, 10^4 perturbed piecewise models were generated by adding a random offset to each of the gain values in table 2.10 and recalculating $\Delta\nu_{\mathrm{noise}}$ with the perturbed model. Each offset was chosen from a uniform distribution between -1 and 1 dB. The mean value for this test was 1.38 GHz with a standard deviation of 0.07 GHz. We therefore quote $\Delta\nu_{\mathrm{noise}} = (1.4 \pm 0.1)$ GHz.

Radiometer reception bandwidth. The radiometer reception bandwidth $\Delta\nu_{\mathrm{rec}}$ is the bandwidth used to characterize the predetection radio-frequency bandpass of the receiver when calculating the sensitivity through the radiometer equation described in section 2.1.3.3. It is defined as

$$\Delta\nu_{\mathrm{rec}} = \frac{\left[\int_0^\infty G(\nu) d\nu\right]^2}{\int_0^\infty G^2(\nu) d\nu}. \tag{2.17}$$

Using the integral result from equation (2.16) and the data in table 2.10, we find

$$\Delta\nu_{\mathrm{rec}} = \frac{\left(\sum\limits_k I_k^{(1)} \right)^2}{\sum\limits_k I_k^{(2)}} = 2.57 \text{ GHz}. \tag{2.18}$$

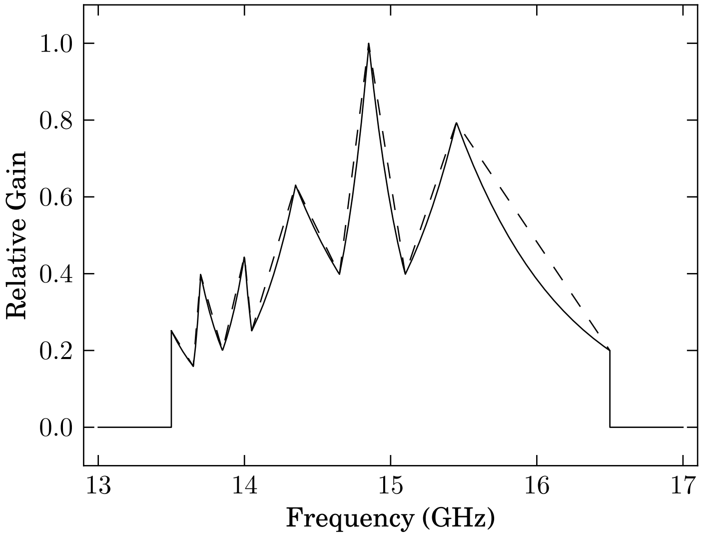

Figure 2.13. Comparison of the linear interpolation (dashed lines) with exponential interpolation (solid lines) in a linear plot of the receiver gain. Values have been normalized to unity at the peak response.

For most radiometers, the reception bandwidth is somewhat larger than the half-power or noise bandwidths, so it is not surprising that $\Delta\nu_{\rm rec}/\Delta\nu \approx 1.7$ (e.g., Evans & McLeish 1977). To evaluate the uncertainty, we use the same random perturbation method described for $\Delta\nu_{\rm noise}$, and find a mean value of 2.57 GHz and a standard deviation of 0.06 GHz. Thus, $\Delta\nu_{\rm rec} = (2.6 \pm 0.1)$ GHz. In section 2.1.3.3, we compare this to the observed thermal noise of the receiver.

2.1.3.3 Sensitivity

A very simple model of a direct detection radiometer, shown in figure 2.14, consists of a radio-frequency amplifier, a detector, and a postdetection filter. As derived in, e.g., Evans & McLeish (1977), the *radiometer equation* relates the $\rm rms$ variation in the output signal, ΔT, to the system temperature at the input, $T_{\rm SYS}$:

$$\frac{\Delta T}{T_{\rm SYS}} = K\left(\frac{2\Delta\nu_{\rm noise}}{\Delta\nu_{\rm rec}}\right)^{1/2}. \tag{2.19}$$

Here $\Delta\nu_{\rm noise}$ is the noise bandwidth of the postdetection filter and $\Delta\nu_{\rm rec}$ is the radiometer reception bandwidth of the amplifier, and the detector has been assumed to be a square-law device. The constant K is a factor that can be used to generalize this result to other receiver architectures. If the postdetection filter is a boxcar integrator of integration time τ, the noise bandwidth $\Delta\nu_{\rm noise} = (2\tau)^{-1}$. This gives the more familiar form

$$\frac{\Delta T}{T_{\rm SYS}} = \frac{K}{\sqrt{\Delta\nu_{\rm rec}\,\tau}}. \tag{2.20}$$

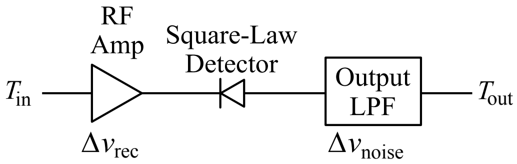

Figure 2.14. Simple radiometer model.

The radiometer equation quantifies the *minimum* achievable noise level with a given receiver and integration time. In practice, additional sources of noise will result in a a higher level. For example, if the gain of the receiver varies significantly over the integration time, output changes due to gain fluctuations will be indistinguishable from changes due to input variation. As a result, gain fluctuations reduce the sensitivity of the receiver. If such fluctuations have rms amplitude ΔG, as shown in Rohlfs & Wilson (2000) this gives

$$\frac{\Delta T}{T_{\mathrm{SYS}}} = \sqrt{\frac{K^2}{\Delta\nu_{\mathrm{rec}}\,\tau} + \left(\frac{\Delta G}{G}\right)^2}. \tag{2.21}$$

Other noise sources will similarly add in quadrature on the right-hand side.

Equation (2.20) with $K = 1$ is valid for a total power radiometer. For a Dicke-switched receiver, $K = 2$. This can be understood as two factors of $\sqrt{2}$: one because only half the integration time τ is spent on-source and another because the source and reference integrations are subtracted, combining their independent noises in quadrature. It is important to note that τ is taken to be the full integration time including *both* Dicke switch states. Although Dicke switching appears to increase the noise level by a factor of two through the K factor, in practice it usually greatly reduces the radiometer noise level by eliminating much of the $\Delta G/G$ factor in equation (2.21).

We now compare the observed noise level to that expected from the radiometer equation. The simplest comparison results from observations of blank sky at zenith, where the input system temperature should be its most stable. In figure 2.15, we plot the first two seconds of the radiometer output measured in DU for the individual beams and the difference between the two beams. These data were collected when the telescope was pointed at zenith for a one-hour period on 17 September 2011.

Using the *ant* and *ref* beam data separately, the receiver acts as a total power radiometer ($K = 1$) with $\tau = 1$ ms per sample. If we compute the difference *ant* − *ref*, the output corresponds instead to a Dicke switched receiver ($K = 2$) with $\tau = 2$ ms per sample. In table 2.11, we use the results from the hour-long data set to estimate the per-sample rms noise and compare this to the results of the radiometer equation. For convenience, we measure T_{SYS} and ΔT in DU rather than converting to K. In this case, T_{SYS} is the average level of the radiometer input—for the switched case, this is the average of the level for the two beams.

Both the *ant* and *ref* data give consistent results. This demonstrates that the two signal paths have nearly equal T_{SYS} and bandpass contributions from the sections of the signal path not in common, i.e., skyward of

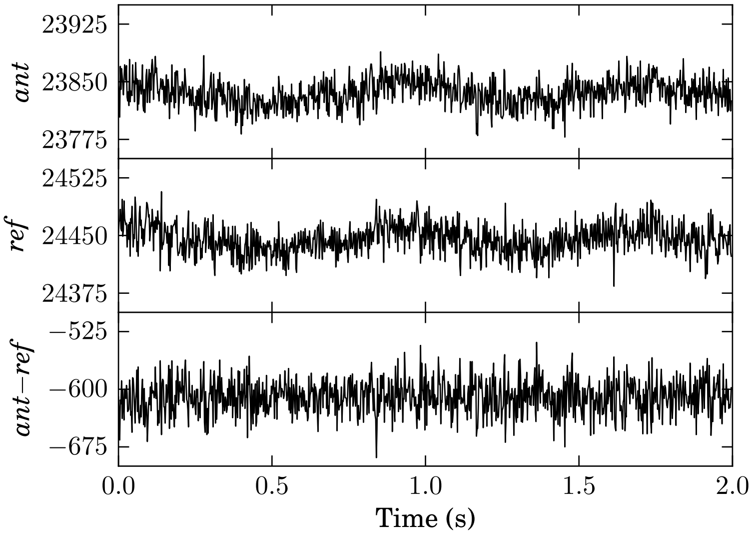

Figure 2.15. Uncalibrated full-rate (500 Hz) Dicke-switched radiometer samples (in DU) collected while pointed at zenith on 17 September 2010. The differenced signal is much flatter, indicating a reduction in slow fluctuations due to atmospheric or receiver variations or other common-mode noise sources. This is at the cost of a $\sqrt{2}$ increase in the white noise level, visible as an increase in the high-frequency scatter in the bottom panel.

Table 2.11. Calculation of receiver sensitivity and comparison with the radiometer equation

Input	K	τ (ms)	T_{SYS} (DU)	ΔT (DU)	Expected ΔT (DU)	$\Delta G/G$
ant	1	1	23,844.2	24.6	14.8	8.2×10^{-4}
ref	1	1	24,453.4	25.1	15.2	8.2×10^{-4}
ant − *ref*	2	2	24,148.8	22.9	21.2	3.6×10^{-4}

Note: Summary of data and calculations comparing the noise in one hour of blank-sky data (the first two seconds are shown in figure 2.15) to the radiometer equation. $K = 1$ for the total power radiometer mode and $K = 2$ for the Dicke-switched mode. Expected ΔT is computed from equation (2.20) and $\Delta G/G$ is computed from equation (2.21) assuming any excess noise results from gain fluctuations.

the Dicke switch. Thus, the difference in DC signal levels between the *ant* and *ref* signals is mostly due to a gain mismatch between the two signal paths rather than a source of excess noise in one.

The switched data match the radiometer equation more closely than do those for the individual channels. In fact, although the predicted ΔT is greater for the Dicke switched mode (21.2 DU versus 14.8 DU or 15.2 DU), the switching eliminates enough excess noise due to atmospheric and receiver gain fluctuations that the measured rms is lower (22.9 DU versus 24.6 DU or 25.1 DU). Trading a factor of two increase in the theoretical noise level (i.e., letting $K = 2$ in equation (2.21)) for the elimination of gain fluctuations through Dicke switching actually lowered the measured noise level.

2.1.3.4 Gain Fluctuations and *1/f* Noise

So far in the sensitivity discussion, we have implicitly assumed that fluctuations are spectrally white—that they are constant in amplitude over all frequencies. This leads to the result, embodied in the radiometer equation (equation (2.20)) that increasing the integration time will reduce the uncertainty in the measurement. This is generally a reasonable assumption for short times, but often breaks down for long integrations. The reason is the presence of noise processes with amplitudes that increase at lower frequencies. These processes are characterized by a power spectral density, $\Phi(f) \propto 1/f^{\alpha}$ for some $\alpha \approx 1$, and are frequently referred to as "red," "pink," or simply "$1/f$" noise. Such noise processes are ubiquitous in nature, and have been observed in systems ranging from turbulence scale distributions in lakes to gain fluctuations in semiconductor devices such as amplifiers and detector diodes (e.g., Schmid 2007; van der Ziel 1988). These latter phenomena affect radiometer sensitivity through the gain fluctuation term in equation (2.21). Note that in this section, we refer to frequency as f rather than ν, both to match the literature on this topic, and to avoid confusion between radio frequency, ν and frequencies in the postdetection signal, f.

On 27 October 2010, 27 min of data were collected while the telescope was pointed at blank sky. Figure 2.16 shows the power spectral density of the average of the *ant* and *ref* signals during 10 min of this period. The $1/f$ behavior is evident, as is contamination due to mains power at 60 Hz and harmonics and a signal from the cryogenic compressor cycling at about 1 Hz and harmonics. The white noise limit at high frequencies is estimated from data, neglecting the narrowband contamination, to be $0.53\ \mathrm{DU}^2\,\mathrm{Hz}^{-1}$. The "knee frequency," f_{knee}, is the frequency at which the $1/f$ noise component equals the white noise, leading to a total noise that is double the white noise floor. This occurs at about $f_{\mathrm{knee}} = 17$ Hz.

In figure 2.17, we show the power spectral density for the differenced data, *ant* − *ref*, during the same 10 min period. The white noise level is found to be $2.07\ \mathrm{DU}^2\,\mathrm{Hz}^{-1}$, or a factor of about four higher, as expected when comparing the variance of a difference to that of an average. There is no evidence for the onset of $1/f$ noise in this plot—it has been reduced tremendously by the differencing. Although in the ideal case *all* $1/f$ noise due to sources in the common signal path would be eliminated, in practice there is some residual effect due to imbalances between the inputs in the two switch states.

In figure 2.18, we plot the power spectral density of the difference *ant* − *ref* over the entire 27 min period, computed to a much lower minimum frequency. Because this increases the number of data points in the plot enormously, we have downsampled by a factor of 30 to reduce the number of points. The white noise level from these data is found to be $2.10\ \mathrm{DU}^2\,\mathrm{Hz}^{-1}$ in agreement with figure 2.17. It appears that residual $1/f$ noise is becoming significant at very low frequencies, with an estimated $f_{\mathrm{knee}} = 5\ \mathrm{mHz}$.

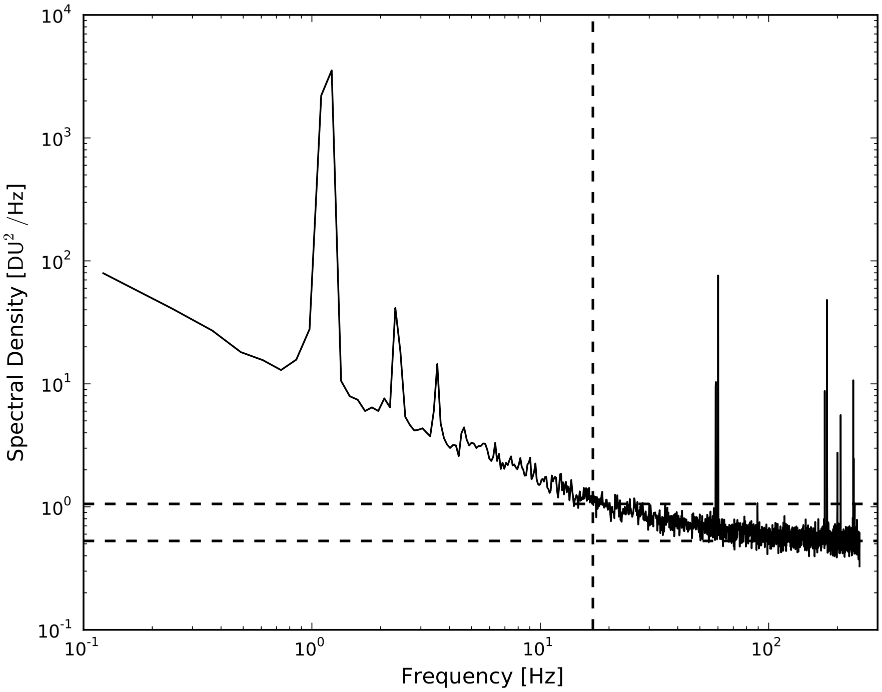

Figure 2.16. Power spectral density of 10 min of averaged *ant* and *ref* samples, illustrating clear $1/f$-type behavior at low frequencies. The horizontal lines indicate $1\times$ and $2\times$ the white noise level ($0.53\ \mathrm{DU}^2\ \mathrm{Hz}^{-1}$). The vertical line indicates the estimated knee frequency, $f_{\mathrm{knee}} \approx 17$ Hz.

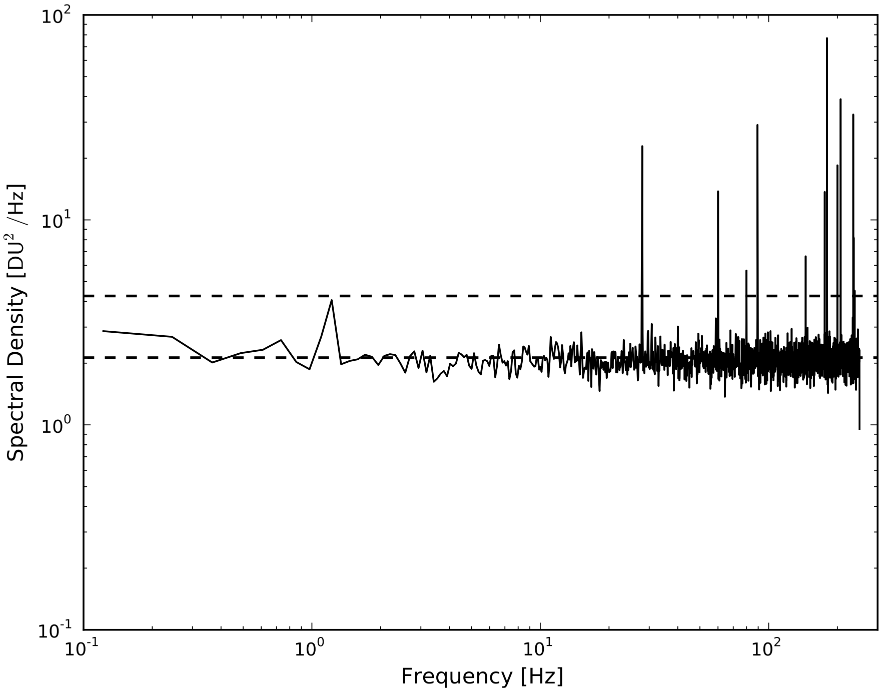

Figure 2.17. Power spectral density of the same 10 min of data shown in figure 2.16, now computed from the difference signal *ant* − *ref*. The reduction in $1/f$ noise is evident. The horizontal lines indicate $1\times$ and $2\times$ the white noise level ($2.07\ \mathrm{DU}^2\ \mathrm{Hz}^{-1}$).

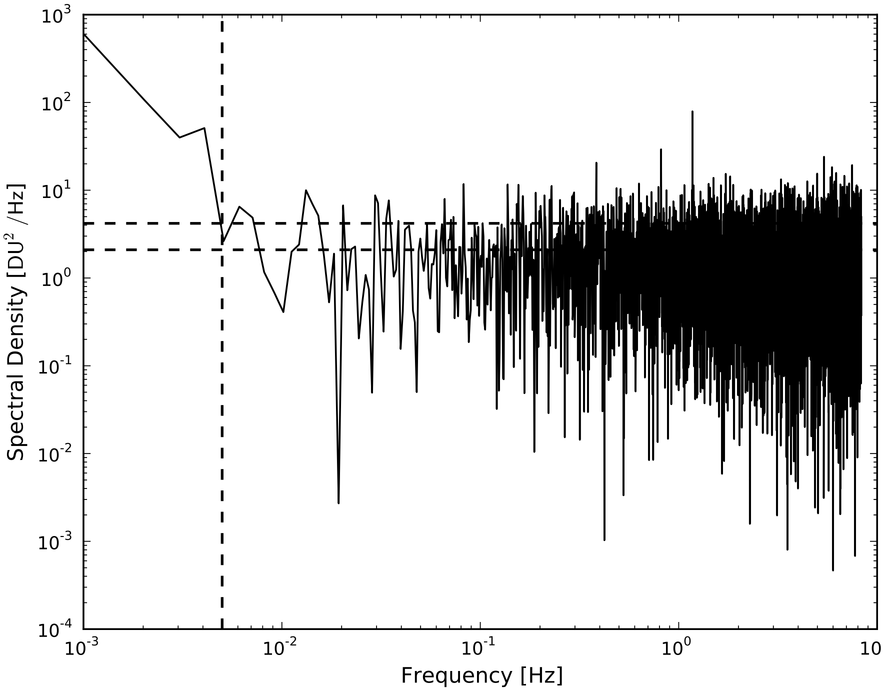

Figure 2.18. Power spectral density of about 27 min of data that overlaps the 10 min interval shown in figures 2.16 and 2.17. The spectrum is computed with finer resolution to accurately measure very low frequencies. To keep the number of data points manageable, the data were down-sampled by a factor of 30 and averaged. The horizontal lines indicate $1\times$ and $2\times$ the white noise level ($2.10\ \mathrm{DU}^2\ \mathrm{Hz}^{-1}$). The vertical line indicates the estimated knee frequency, $f_{\mathrm{knee}} \approx 5\ \mathrm{mHz}$. The "bushy" appearance of the plot at high frequencies is a visual artifact of the logarithmic binning—there is no actual increase in the white noise level at high frequencies.

The radiometer equation can be used to estimate the reception bandwidth of a receiver from the white noise level in the power spectral density of the output. Given a white noise power spectral density Φ_0 in units of power per Hz, we use $\Delta\nu_{\rm noise} = 1$ Hz in equation (2.19) and find

$$\Delta\nu_{\rm rec} = \frac{2\,K^2\,{T_{\rm SYS}}^2}{\Phi_0}. \tag{2.22}$$

In place of $T_{\rm SYS}$, we insert the average level of the input signal in DU. The average value of the input during this test was 2.5265×10^4 DU. We find $\Delta\nu_{\rm rec} = 2.41$ GHz using the averaged data ($K = 1$ and $\Phi_0 = 0.53\ {\rm DU}^2\ {\rm Hz}^{-1}$) and $\Delta\nu_{\rm rec} = 2.47$ GHz from the differenced data ($K = 2$ and $\Phi_0 = 2.07\ {\rm DU}^2\ {\rm Hz}^{-1}$). This agrees reasonably well with the value $\Delta\nu_{\rm rec} = (2.6 \pm 0.1)$ GHz we computed from the receiver bandpass in section 2.1.3.2.

2.1.3.5 Calibration Diodes

A pair of calibrated noise diodes, referred to as the NOISE and CAL diodes, are connected to the main beam input via directional couplers to the Dicke switch. At their outputs, these noise diodes provide an excess noise ratio of (31 ± 1) dB from 12–18 GHz with compensation to maintain output stability with temperature. The outputs of the noise diodes are reduced in amplitude without introducing excessive thermal noise by connecting them to the *ant* signal chain through directional couplers as shown in figure 2.10.

These calibration diodes provide two equivalent noise temperatures for calibration. The NOISE diode provides a noise temperature comparable to the system temperature and the CAL diode provides a noise temperature comparable to the antenna temperature of the astronomical sources we are observing. The equivalent noise temperatures of the NOISE and CAL diodes at the receiver input are about 67.3 K and 1 K—see figures 2.23 and 2.25.

The temperature stability of the NOISE and CAL diodes was measured using a calibrated continuous-wave RF power meter to measure the output level as the temperature was raised from room temperature (near 300 K) to about 325 K using a hot air gun. The diode under test was removed from the receiver and its metal case was bolted to a thick aluminum plate. The hot air gun was applied to the back side of the plate away from the diode and a few seconds were allowed for the diode to equilibrate with its case temperature before the output of the power meter was recorded. The case temperature was measured using an infrared thermometer (specified accuracy ± 0.2 K) aimed at a piece of tape with high infrared emissivity that was securely attached to the diode case. After the temperature was raised to the maximum tested, the diode was allowed to cool back to room temperature. Measurements were made both during heating and cooling.

Both diodes were tested, but the more accurate absolute power reading was only recorded for the CAL diode. The results are plotted in figure 2.19, showing variation of about $8 \times 10^{-4}\ {\rm dB} \cdot {\rm K}^{-1}$ assuming a linear model. The NOISE diode exhibited nearly identical relative output measurements, so we believe its temperature stability to be similar.

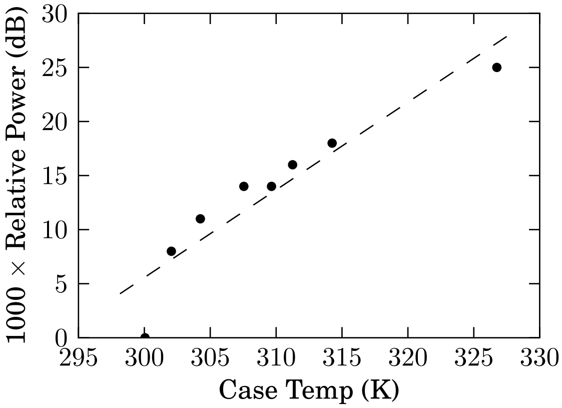

Figure 2.19. Change in absolute CAL diode power output versus its case temperature. The dashed line is a linear fit with slope 8×10^{-4} $\mathrm{dB} \cdot \mathrm{K}^{-1}$.

Table 2.12. Results of a stability test of CAL diode output versus diode case temperature

Case Temp (°C)	Relative (dB)	Absolute (dBm)
26.9	0.00	−40.457
36.5	0.01	−40.443
41.1	0.02	−40.439
53.6	0.02	−40.432
38.1	0.01	−40.441
34.4	0.01	−40.443
31.1	0.01	−40.446
28.9	0.01	−40.449

Table 2.13. Results of a stability test of NOISE diode output versus diode case temperature

Case Temp (°C)	Relative (dB)
27.1	0.00
28.6	0.00
31.7	0.00
35.6	0.01
36.1	0.01
39.7	0.01
44.4	0.01
48.4	0.02
45.0	0.02
33.7	0.00
28.8	−0.01

Note: During this test, the absolute output was not recorded.

Although the calibration diodes' physical temperatures are not regulated or monitored directly, they are located inside the warm section of the receiver enclosure. This is regulated at an air temperature of approximately 300 K with rms fluctuations of about 1.1 K. In 2009, during hot weather when the ambient air temperature reached $\gtrsim 310\ K$, excursions up to about 307 K inside the enclosure were observed. During winter months occasional brief decreases as low as 288 K occurred. Using the temperature coefficient we measured, this corresponds to a $0.015\ \mathrm{dB}$ (0.3%) change in the output level from one extreme to the other if the diodes vary over the full range of temperatures in the enclosure. If we fit a quadratic instead of linear model to the data in figure 2.19, we predict a $0.031\ \mathrm{dB}$ (0.7%), still well less than 1%.

Even this level of temperature-induced output variation is unlikely to occur on short timescales. Temperature swings this large occur between seasons, not diurnally. We therefore need not worry about variations in calibration diode output due to temperature changes in the course of a day—we may treat their outputs as constant to the level of precision we require.

2.1.3.6 Beam Isolation

The Dicke switch is a ferrite switch that couples one of two feed-side ports to the common port connected to the radiometer. The device is controlled electronically by reversing the magnetic field domains in a ferrite. Such an electronic device is the only practical microwave RF switch capable of continuous 500 Hz switching for years. A trade-off must be made in the isolation between the *ant* and *ref* ports—the isolation of a ferrite switch is much lower than a mechanical switch, but the latter would wear out in a very short time. Furthermore, devices of this sort designed for cryogenic operation were not available at the time the receiver was constructed (and are now extremely rare in any form). Thus, the Dicke switch is a device intended for room temperature operation that was found to behave reasonably at temperatures down to about 70 K (Leitch 1998). As the temperature is lowered below this, the isolation drops to unacceptable levels.

By observing the response in the *ref* channel when the NOISE diode (in the *ant* channel), the isolation of the Dicke switch is found to be about $15\ \mathrm{dB}$. The isolation need not be symmetric, but Leitch (1998) found the two directions to match to about $0.5\ \mathrm{dB}$. Because the isolation is nearly symmetric, the effect of isolation is to reduce the effective gain for a flux density measurement measured as described in section 2.2.2.2. Because we observe our astronomical calibration sources with the same procedures as our program sources, therefore, an explicit isolation correction is not required.

2.1.3.7 Nonlinearity

Although amplifiers and detectors are designed to exhibit linear response, deviation from linearity at some level is inevitable. The most important nonlinear effect for our receiver is gain compression, a reduction in gain as the output level increases. This effect is demonstrated in figure 2.20, where the difference between the ideal linear response (solid line) and the actual response (dashed line) is due to compression.

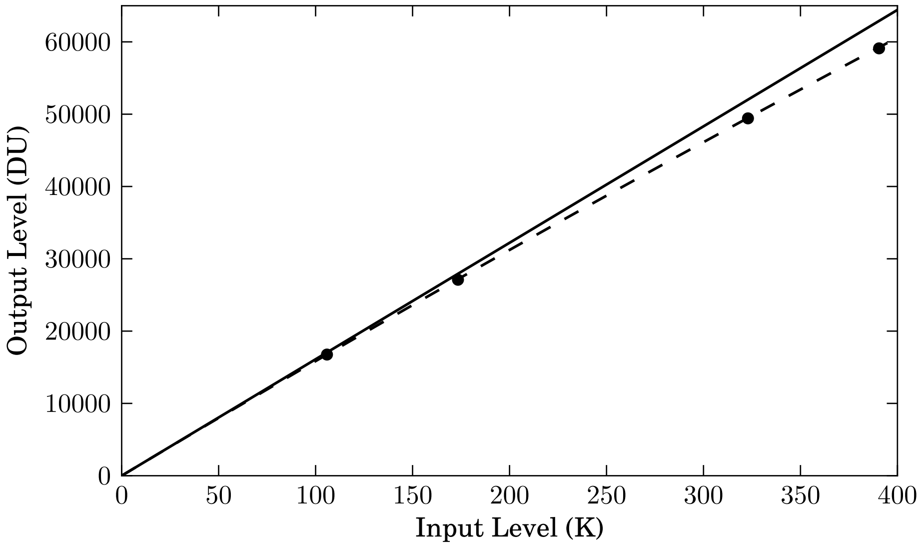

Figure 2.20. One trial of a hot/cold Y-factor measurement demonstrates the presence of gain compression in the receiver. The four input levels, here measured in K, correspond to the LN2 cold load (assumed to be 77.4 K) and the ambient-temperature hot load (measured to be 294.4 K), each with and without the NOISE diode enabled (equivalent noise temperature measured to be 67.6 K). The measured receiver temperature from this trial, 28.5 K, is also included. The solid line is a constant gain of 161 $\mathrm{DU} \cdot \mathrm{K}^{-1}$ with no compression. The dashed line through the data points is the same gain, but including compression with $b = -9.824 \times 10^{-7}\ \mathrm{DU}^{-1}$.

Modeling the exact nonlinear behavior of a receiver system is complicated, but for our needs, a simple empirical model is sufficient. We use the same nonlinearity model developed in Herbig (1994) and also described in Leitch (1998), where deviation from linear response is characterized by a single parameter, b. The model simply adds a gain factor proportional to the output level to the constant component of the gain,

$$y = G\,(1 + b\,y)\,x, \tag{2.23}$$

where x and y are the input and output signals, respectively, and G is the nominal gain measured well below where compression effects are significant. Here, b is given in inverse units of the output signal. This model is reasonable for responses that exhibit slight compression, when $|b| \ll y^{-1}$. Figure 2.21 demonstrates the effect of various b values. For a system with gain compression, $b < 0$. When $b > 0$, the model response describes a gain expansion.

Equation (2.23) describes the response of the system for any input. For many measurements, we are interested in the small-signal gain. To find this, we solve for y as a function of x, giving

$$y = \frac{G\,x}{1 - b\,G^{-1}\,x}. \tag{2.24}$$

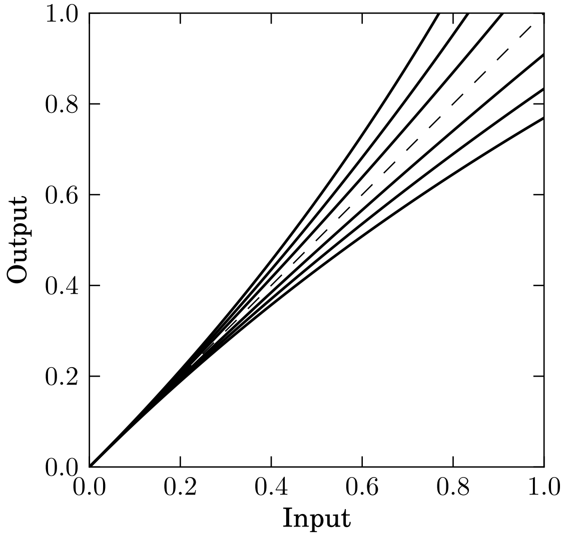

Figure 2.21. Behavior of the compression model given by equation (2.23) for $b = 0$ (dashed line) and -0.3, -0.2, -0.1, 0.1, 0.2, 0.3 (solid lines, bottom to top).

The small-signal gain near an input operating point x_0 with corresponding output y_0 is then given by the derivative at that point,

$$\left.\frac{dy}{dx}\right|_{x_0} = \frac{G}{\left(1 - b\,G^{-1}\,x_0\right)^2} = G\,\left(1 + b\,y_0\right)^2\,. \tag{2.25}$$

This is the correction that would apply to a differential measurement, such as when Dicke switching is employed.

This compression model is convenient because the correction can be applied given the output level and a measured value for b. For a measured output y, the corrected input level is simply

$$x = \frac{G^{-1}\,y}{1 + b\,y}. \tag{2.26}$$

For a differential measurement Δy, the input difference is given by

$$\Delta x = \frac{G^{-1}\,\Delta y}{\left(1 + b\,y\right)^2}. \tag{2.27}$$

Our receiver is constructed of several components that each individually exhibit compression. As described in Leitch (1998), the total compression of the system will depend on the signal level presented to each component and will change if the attenuation between components changes. Most significantly, the nonlinearity parameter must be measured for each setting of the programmable attenuator shown in figure 2.10. During normal observations, the attenuator is very rarely adjusted, so we simply measure the overall b parameter for the operating attenuation. However, during hot/cold calibration procedures, a higher attenuation level is used, so b must be measured separately for these procedures.

2.2 Radiometry Techniques

In order to make use of the hardware described in the previous section, its performance characteristics must be measured and understood and measurement techniques must be specified. In this section, we first describe several calibration procedures that are performed to measure and monitor the receiver performance. We then describe the measurement procedures that are used as the basis for the monitoring program.

2.2.1 Calibration and Diagnostic Procedures

Regular diagnostic tests must be performed to ensure that the telescope and receiver are performing as expected and that the data can be properly calibrated. Fortunately, the 40 m telescope and receiver have been very stable and reliable, simplifying these tasks a great deal. Calibrations have been performed regularly, on a schedule that has evolved as we have better understood the stability and requirements of the instrument.

From 2007 until November 2008, receiver calibrations were performed occasionally but without a regular schedule. Beginning in November 2008, approximately monthly calibrations were performed. In December, 2009, the schedule was reduced to semimonthly, which has proven adequate. These calibrations include Y-factor measurements to characterize receiver temperatures, sky dips to measure atmospheric optical depth and to determine the ground spillover, calibration diode effective temperature measurements, and observations of calibration sources to measure the aperture efficiency. We describe these methods here (see section 2.1.2.1 for the aperture efficiency measurement procedure).

2.2.1.1 Measuring Receiver and Calibration Diode Noise Temperatures

We use a single calibration procedure to determine the receiver noise temperature and to find the equivalent noise temperature of the calibration diodes. The basis for our measurement is the Y-factor method (e.g., Evans & McLeish 1977; Rohlfs & Wilson 2000).

To measure the receiver noise temperature using the Y-factor method, two input loads of different known brightness temperatures are needed. The Y-factor is defined as

$$Y \equiv \frac{P_H}{P_C}, \tag{2.28}$$

where P_H and P_C are the receiver output levels when presented with hot and cold loads at temperatures T_H and T_C, respectively. For a receiver with a noise temperature T_{RX} and an input temperature T,

$$P = G\left(T + T_{\mathrm{RX}}\right). \tag{2.29}$$

Using the two loads and solving, we find

$$T_{\mathrm{RX}} = \frac{T_H - Y\,T_C}{Y - 1}. \tag{2.30}$$

The uncertainty in this measurement is then given by

$$\sigma^2_{T_{\mathrm{RX}}} = \left(\frac{1}{Y-1}\right)^2 \sigma^2_{T_H} + \left(\frac{Y}{Y-1}\right)^2 \sigma^2_{T_C} + \left(\frac{T_H - T_C}{(Y-1)^2}\right)^2 \sigma^2_Y, \tag{2.31}$$

where σ_x^2 is the mean-squared uncertainty in the value of x. Note that in equation (2.29), the receiver response was assumed to be linear with no offset. In practice, neither assumption is valid, so corrections must be applied.

The procedure we follow when performing a hot/cold load test comprises a series of five measurements performed with the receiver on the telescope:

1. fill the feed aperture with the hot load and measure P_H;
2. enable the NOISE diode and measure P_{H+N};
3. disable the NOISE diode, disconnect the radiometer input from the DAQ back end and measure P_Z;
4. remove the hot load, fill the feed aperture with the cold load, and measure P_C; and
5. enable the NOISE diode and measure P_{C+N}.

This procedure is repeated several times. The results are averaged and their scatter is used to estimate the uncertainty in the result.

Both the hot and cold loads are constructed from blocks of radio-absorbent foam mounted with a reflective metal backing. The hot load is kept at the ambient temperature and is instrumented with a thermometer to measure T_H. The cold load is enclosed in a radio-transparent insulating box that is filled with liquid nitrogen (LN2), so $T_C \approx 77$ K. Care must be taken to minimize uncertainty in T_H and T_C, e.g., by giving the cold load time to equilibrate with its LN2 bath, by keeping the hot load in the shade to avoid direct solar heating, and by working quickly to minimize changes in ambient temperature and avoid excessive LN2 boil-off. Even with these precautions, it is difficult to precisely estimate the effective load brightness temperature, which depends both on the physical temperature of the LN2 and the effect of reflections due to imperfect matching and absorption by the absorbing foam. Based on experience, the effective 2 cm brightness temperature for a well-matched box load is likely within 1–2 K of the physical temperature (D. P. Woody and J. W. Lamb, personal communication). We adopt a value of $T_C = 77.4$ K for the cold load, which may be a slight underestimate of its true brightness temperature. For the typical value $Y \approx 3$ in our tests, each 1 K underestimate in T_C would result in a $\sim$1.5 K overestimate of T_{RX}.

The hot load brightness temperature is similarly uncertain due to reflection and imperfections in the load, as well as the possibility that the thermometer does not reflect the physical temperature of the load due to thermal gradients in the foam. It is more difficult to estimate the uncertainty because reflections from the load may be terminated on the ground ($T \approx 300$ K), on the sky ($T \approx 10$ K, including the CMB contribution), or within the cryostat. In any case, the temperature of the terminating material is likely to be similar to or less than the measured temperature of the hot load, so the measured T_H is probably a slight overestimate of the true brightness temperature. Each 1 K overestimate in T_H increases the measured T_{RX} by ~ 0.5 K for our typical value $Y \approx 3$. Thus, the receiver temperatures we report are perhaps slightly conservative.

Given a single trial, the data are reduced as follows. First, we estimate the nonlinearity parameter b_{HC} in equation (2.23) applicable for the measurement. Because the ~ 300 K temperature of the hot load dictates a higher programmable attenuator setting for this test than is used for normal observing, this measurement of b_{HC} cannot be used to correct other data. We find b_{HC} by comparing the change in receiver output due to the NOISE diode when measured against the hot load ($\Delta_H \equiv P_{H+N} - P_H$) versus when measured against the cold load ($\Delta_C \equiv P_{C+N} - P_C$). The parameter b_{HC} is given by

$$b_{HC} = \frac{2\left(\sqrt{\Delta_C} - \sqrt{\Delta_H}\right)}{\Sigma_C \sqrt{\Delta_H} - \Sigma_H \sqrt{\Delta_C}}, \tag{2.32}$$

where $\Sigma_H = P_H + P_{H+N} - 2P_Z$ and $\Sigma_C = P_C + P_{C+N} - 2P_Z$ are the average offset-corrected output levels.

Next, using equation (2.26) we correct our measured values for nonlinearity and offset using the measured value for b_{HC}. For P_H, this is

$$P'_H = \frac{P_H - P_Z}{1 + b_{HC}\left(P_H - P_Z\right)}, \tag{2.33}$$

and similar for P_C, etc. Now using the corrected P'_H and P'_C values to compute Y, equation (2.30) yields T_{RX}.

The receiver gain (in $\mathrm{DU} \cdot \mathrm{K}^{-1}$) at this attenuator setting can be calculated as

$$G_{HC} = \frac{P'_H - P'_C}{T_H - T_C}. \tag{2.34}$$

This can be used to convert the measured noise diode response from DU into its equivalent noise temperature in K:

$$T_{\mathrm{NOISE}} = G_{HC}^{-1} \times \left(P'_{H+N} - P'_H\right). \tag{2.35}$$

Although b_{HC} and G_{HC} are only applicable to the attenuator setting used for the hot/cold load test, T_{NOISE} is a property of the NOISE diode output and does not depend on the attenuator. The equivalent temperature of the CAL diode, T_{CAL}, is determined from the ratio of its response to that of the NOISE diode. Because the response of the CAL diode is much smaller than that of the NOISE diode, this comparison is performed at the normal operating attenuator setting rather than at the hot/cold load test setting.

Results from the Y-factor tests and the related measurements are shown in figures 2.22 through 2.27. In figure 2.22, the large T_{RX} outlier in February 2010 resulted from a hardware problem in the receiver that was corrected shortly after the measurement.

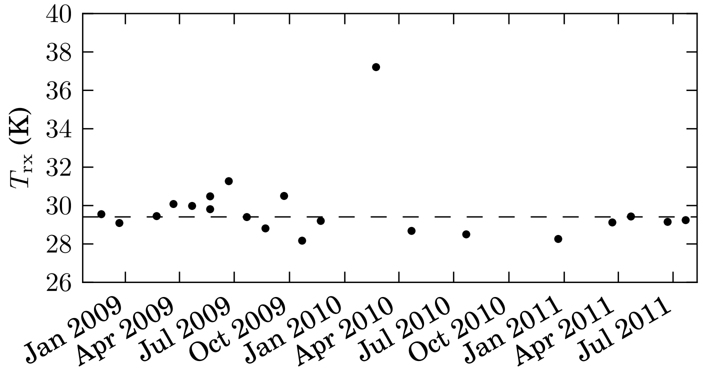

Figure 2.22. Receiver temperature measured between November 2008 and July 2011. Error bars estimated on individual measurements, typically about 0.15 K, are smaller than the plotted points. Dashed line indicates the mean, $T_{\mathrm{RX}} = (29.4 \pm 0.2)$ K, excluding the outlier in February 2010 and using the sample standard deviation to estimate the uncertainty.

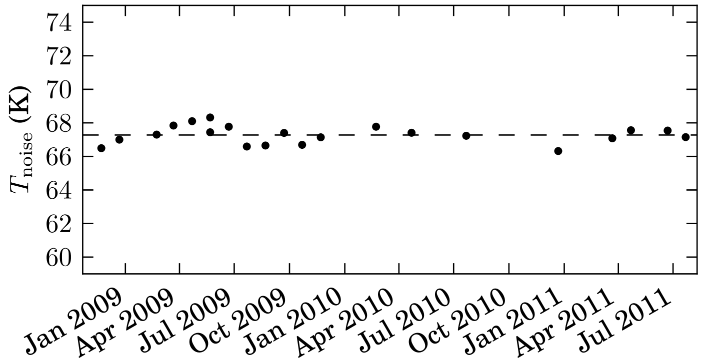

Figure 2.23. NOISE diode equivalent noise temperature measured between November 2008 and July 2011. Error bars estimated on individual measurements, typically about 0.05 K, are smaller than the plotted points. Dashed line indicates the mean, $T_{\mathrm{NOISE}} = (67.3 \pm 0.1)$ K, where the uncertainty is estimated from the sample standard deviation.

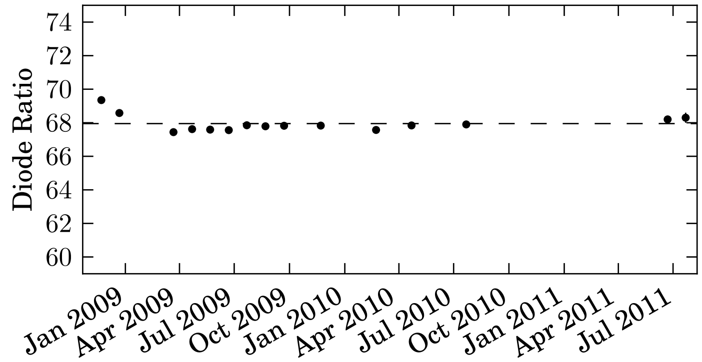

Figure 2.24. Nonlinearity-corrected NOISE to CAL output ratio measured between November 2008 and July 2011. Error bars estimated on individual measurements, typically about 0.06, are mostly smaller than the data points. Dashed line indicates the mean ratio, (68.0 ± 0.1), where the uncertainty is estimated from the sample standard deviation.

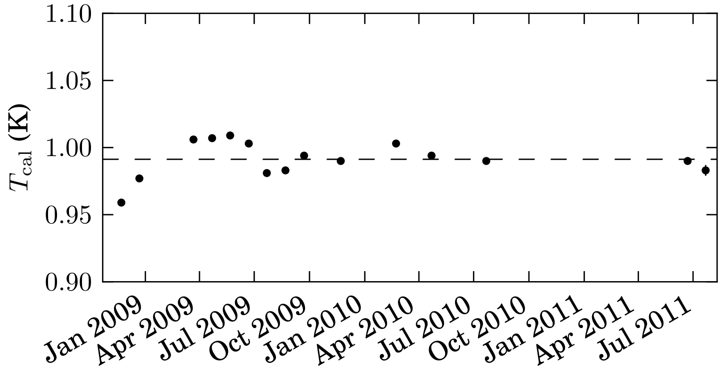

Figure 2.25. CAL diode equivalent noise temperature measured between November 2008 and July 2011. Error bars estimated on individual measurements, typically about 0.001 K, are mostly smaller than the data points. Dashed line indicates the mean, $T_{\rm CAL} = (0.991 \pm 0.003)$ K, where the uncertainty is estimated from the sample standard deviation.

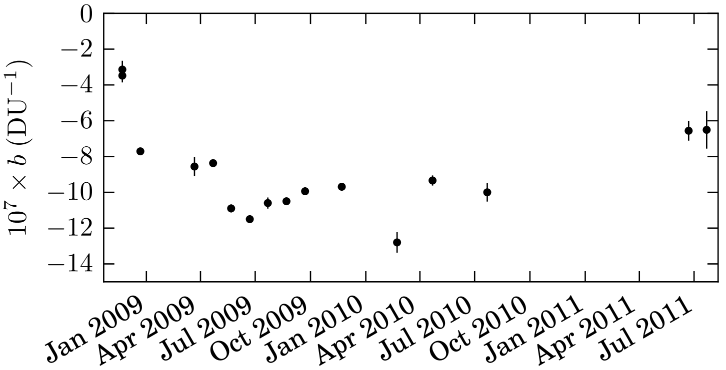

Figure 2.26. Nonlinearity parameter b measured during skydip tests between November 2008 and July 2011. The uncertainty is typically about 10^{-8} and many error bars are smaller than the data points. The 2008 data were measured with the programmable attenuator set to 5, leading to a smaller $|b|$ on those dates. The decrease in $|b|$ in the June and July 2011 may result from a change in receiver gain following maintenance or from a change in the measurement procedure.

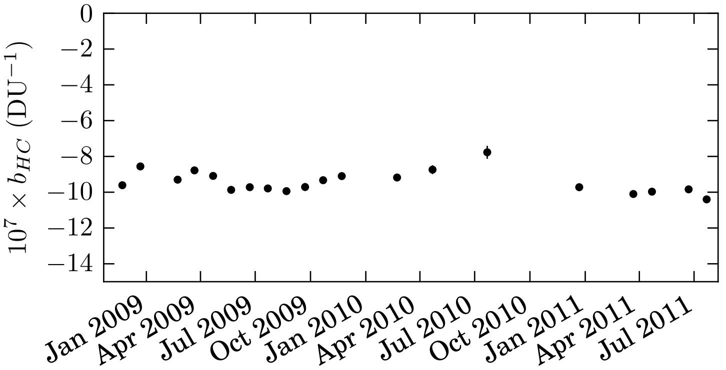

Figure 2.27. Nonlinearity parameter $b_{\rm HC}$ measured during hot/cold load tests between November 2008 and July 2011. The uncertainty is typically about 10^{-8} and many error bars are smaller than the data points.

2.2.1.2 Measuring Atmospheric and Ground Pickup

In order to measure the emission due to the atmosphere, we perform a "sky dip" measurement. That is, we measure the receiver output as we dip the telescope from near zenith to near the horizon. At lower elevations, the path length through the atmosphere—the air mass—increases. The air mass at the moderate elevations at which we observe is well approximated by assuming a plane-parallel atmosphere, where the air mass d_{atm} at zenith angle z in atmospheric depths is

$$d_{\text{atm}} \approx \sec z. \tag{2.36}$$

If we assume that the absorption coefficient α of the atmosphere is uniform, then the atmospheric optical depth is simply

$$\tau = \int_0^{d_{\text{atm}}} \alpha \, dy = \alpha \, d_{\text{atm}}. \tag{2.37}$$

When observing an astronomical source with a brightness temperature T_s, the detected signal will then be

$$\begin{aligned} T &= (T_s + T_{\text{CMB}}) \, e^{-\tau} + T_{\text{ATM}} \left(1 - e^{-\tau}\right) + T_{\text{GND}} + T_{\text{RX}} && (2.38) \\ &\approx (T_s + T_{\text{CMB}}) \, (1 - \alpha \, d_{\text{atm}}) + T_{\text{ATM}} \, \alpha \, d_{\text{atm}} + T_{\text{GND}} + T_{\text{RX}}, && (2.39) \end{aligned}$$

where T_{ATM} is the atmospheric temperature, assumed constant, T_{GND} is the temperature due to ground pickup, and T_{RX} is the receiver noise temperature. The latter approximation applies if $\tau \ll 1$. We are normally justified in assuming T_{GND} is constant over short times because the antenna sidelobes that couple it to the receiver are extended in solid angle, so it measures the average temperature over a large area of the ground. Exceptions to this occur, for example, during periods when the ambient temperature is changing rapidly, when localized radio-frequency interference sources are present, or when the telescope is tipped to low enough elevation that near-in sidelobes intersect the ground. The latter can occur at relatively high elevations when pointed toward nearby mountains, but as these are less than 10° elevation, this is not a concern for normal observing.

Figure 2.28 shows the results for a single sky dip performed on 27 April, 2011. The radiometer data were corrected for nonlinearity, then converted to temperatures by comparing their response to NOISE diode measurements just before the sky dip. In panel (c), the dashed line is the fit

$$T_{\text{ant}} = (5.2 \, d_{\text{atm}} + 47.3) \ \text{K}. \tag{2.40}$$

Comparing this to equation (2.39) with $T_s = 0$, we conclude $\alpha \, (T_{\text{ATM}} - T_{\text{CMB}}) = 5.2$ K and $T_{\text{GND}} + T_{\text{RX}} + T_{\text{CMB}} = 47.3$ K. Although the temperature of the atmosphere varies, in the troposphere where most of the water vapor responsible for atmospheric radio absorption and emission is located, temperatures are typically

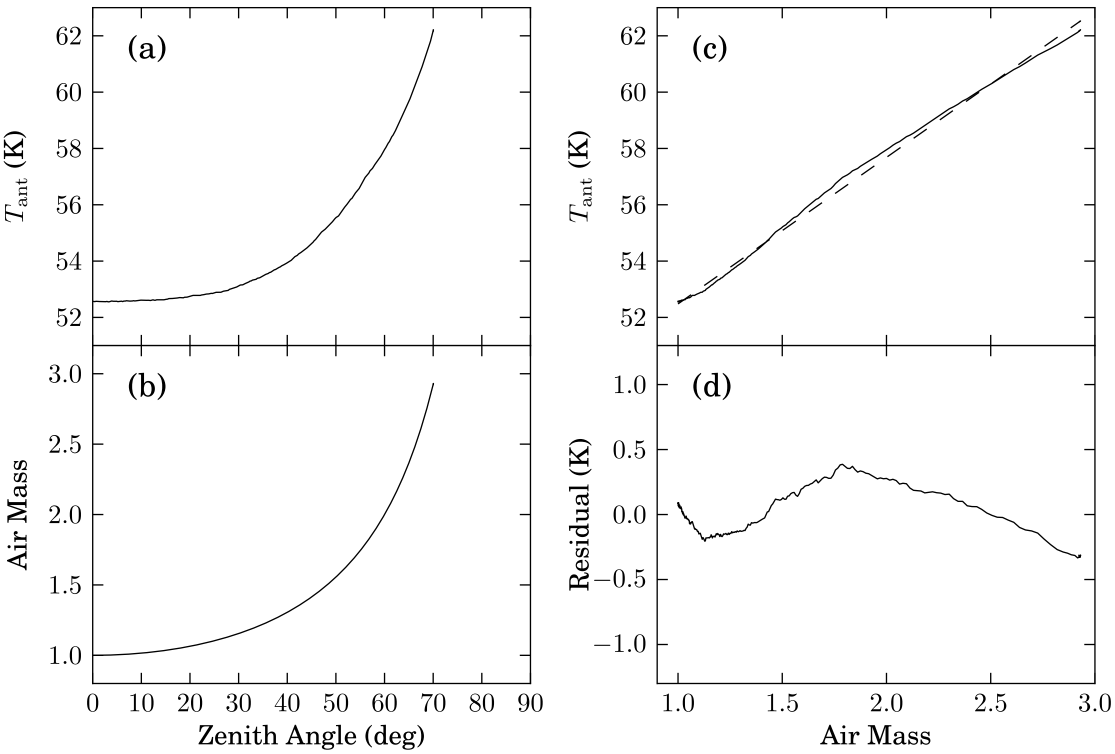

Figure 2.28. Downward portion of a sky dip (27 April, 2011). (a) Calibrated radiometer output as a function of zenith angle. (b) Approximate air mass as a function of zenith angle, computed from equation (2.36). (c) Radiometer output, now plotted against air mass. The dashed line is a linear fit to the data with slope 5.2 K and intercept 47.3 K. (d) Residual after removing the linear fit from the data, showing small nonrandom deviation from the linear model.

230–300 K (e.g, NOAA 2010)), so 270 K is a reasonable estimate of $T_{\rm ATM}$. This gives $\tau_0 \approx 0.02$ as the atmospheric optical depth at zenith, justifying the small-τ approximation in equation (2.39).

The residuals from the fit, shown in panel (d) of figure 2.28, systematically deviate from the linear model—much of this structure is consistent from one sky dip to the next. This most likely results from changes in the ground pickup term as the telescope slews. In particular, the sharp increase near zenith occurs when the telescope feed is looking into the ground, perhaps coupling around the edge of the dish or detecting emission *through* the dish, much of which is perforated rather than solid.

A sky dip procedure is also used to measure the nonlinearity parameter, b, of the receiver, defined by equation (2.23). The calibration diodes produce a constant input noise temperature increment. By measuring the resulting receiver output increment at several background levels, and fitting to the nonlinearity model, b can be estimated. By measuring the calibration diode response at several elevations, the varying atmospheric signal that results from variations in τ with elevation provides this varying background level. Normally the NOISE diode, with a temperature increment of about 70 K, is used for this measurement. The fainter CAL diode signal (~ 1 K) is not strong enough for a reliable measurement of the nonlinearity effect for typical background variations.

2.2.2 Observation Procedures

2.2.2.1 Pointing

Pointing the 40 m telescope accurately is a crucial requirement for this observing program. As discussed in section 2.2.2.2, flux densities are measured using an on-off sky switching, which requires that the 40 m telescope be able to reliably position the center of the beam pattern on the target. Although the beam width of $\theta_{\rm FWHM} \approx 2\overset{\prime}{.}6$ is relatively broad, even a $15''$ pointing error will result in a 2.5% reduction in gain.

Pointing model. The 40 m telescope is equipped with encoders on the azimuth and elevation shafts. These readings and the readings from the tilt sensors (see section 2.1.1.3) are combined in a pointing model that generates encoder azimuth and zenith angle offsets based on the requested position on the sky. The pointing model has 9 terms for the azimuth angle correction and 5 terms for the zenith angle correction,

$$\begin{aligned} \Delta\phi_{model} &= A_1 \sin\theta + A_2 + A_3 \sin\phi \cos\theta \\ &+ A_4 \cos\phi \cos\theta + A_5 \cos\theta + A_6 \sin\phi \sin\theta \\ &+ A_7 \cos\phi \sin\theta + A_8 \sin(4\theta) + A_9 T_{LR} \cos\theta, \end{aligned} \tag{2.41}$$

$$\Delta\theta_{model} = Z_1 + Z_2 \sin\theta - Z_3 \cos\phi + Z_4 \sin\phi + Z_5 T_{AF}. \tag{2.42}$$

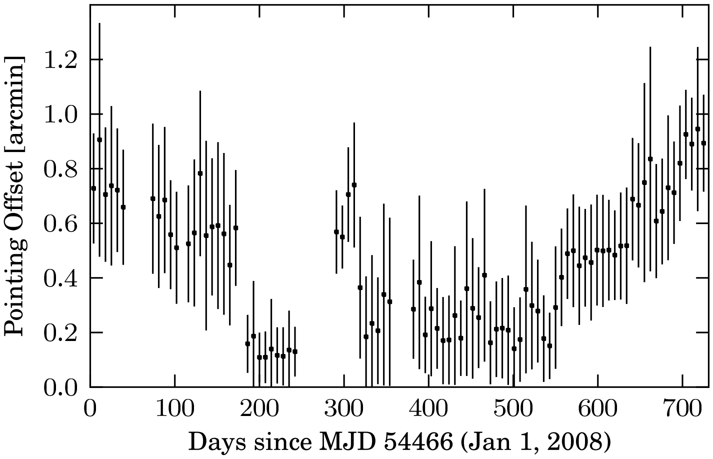

Figure 2.29. Residual error between the pointing model and the actual requested position, plotted in week-long bins for 2008 and 2009. The plotted data and errors are the weekly means and standard deviations of the pointing offsets measured by the pointing calibrations.

Here, ϕ and θ are the requested azimuth and zenith angles, $\Delta\phi_{model}$ and $\Delta\theta_{model}$ are the pointing model corrections for the azimuth and zenith angles, A_i and Z_i are the pointing model coefficients, and T_{AF} and T_{LR} are the aft-forward and left-right tilt meter readings.

The pointing model is determined by performing POINT procedures, described below, at a range of mount coordinates that spans its range. Each POINT produces a measurement of the offset between the encoder readings and the actual position of the telescope beam on the sky. The pointing model coefficients are fit using a least-squares procedure to minimize the error between the model prediction and the measured offsets. The rms residual between the pointing model and the offsets is typically about $15''$ on the sky. This quantifies our ability to point the telescope blindly at a desired position.

We have found that the pointing model terms drift slowly with time. Figure 2.29 shows the residual offset between the pointing model and the actual requested position for 2008 and 2009. The sharp steps in the average offset correspond to adjustments in the pointing model. Before 2010, we adjusted the pointing model two to three times per year to minimize the scatter in the offset and maintain an average offset less than about $0\rlap{.}'5$ to ensure accurate pointing. The data for the fit were collected in a two day period when the telescope was dedicated to measuring the pointing model.

In 2010, we changed our procedure and eliminated the dedicated pointing model measurements. Instead, the POINT procedures executed as part of the monitoring program were collected and the residuals were used to fit pointing model updates as necessary. This eliminates the engineering time overhead of the dedicated measurements and ensures that the model is fit to data collected at mount positions corresponding to those spanned by our program.

Pointing offsets. In addition to the pointing model correction, at least once per hour we measure the pointing offset between a bright pointing calibrator and the model prediction. This measures the effect of wind and thermal loading. In early 2009, we determined that these pointing offsets have the accuracy we require only at separations up to about 30° from the position where the pointing offset was measured. Because of this effect, after MJD 54906 (16 March 2009), care was taken when scheduling to ensure that flux density measurements were always made at separations of less than 15° from the pointing offset measurement. Prior to this, no such limit was in place. We have discarded flux densities measured with a separation of more than 30°.

POINT procedure. Before December 2010, pointing offsets were measured using the POINT procedure, which was implemented in both the VAX and MCS control systems. This procedure operated by performing 3-point cross-scans of the calibrator in both azimuth and zenith angle and fitting a fixed-width Gaussian beam profile to each axis to determine the position of the peak. A pointing offset measurement is considered invalid if its signal-to-noise ratio is less than two, or if the offset indicates that the peak was outside the span of the cross-scan, $\pm\theta_{\rm FWHM}/2$. When using the POINT procedure, several iterations are normally attempted, moving the cross-scan center by up to $\theta_{\rm FWHM}/2$ after each attempt. This allows offsets less than the $\theta_{\rm FWHM}$ to be measured reliably.

POINT2D procedure. In December 2010, the POINT2D procedure was introduced in the MCS control system as a more efficient procedure for measuring pointing offsets. As suggested by its name, the POINT2D procedure measures the two dimensions of the sky offset at once, rather than first scanning in azimuth and then in elevation. In this procedure, measurements are performed at the expected peak position of the beam and 12 other positions in a pair of hexagons centered on the expected peak at $0.4\,\theta_{\rm FWHM}$ and $0.7\,\theta_{\rm FWHM}$ radii. A final point at an azimuth offset about $8'$ from the expected peak is measured to sample the sky background. The pattern of measurement positions is shown in figure 2.30. Using a least-squares fit to the results, the offsets in both directions can be measured simultaneously. This procedure has been found to be more robust and, importantly, faster than the original POINT procedure.

2.2.2.2 Beam Switching and Flux Density Measurements

While Dicke switching does much to reduce the large error terms due to the atmosphere, the ground, and gain fluctuations in the receiver, it does not remove linear drifts in any of these quantities and the situation can be further improved by beam switching. Beam switching in azimuth is optimum because by maintaining a constant elevation we minimize changes to the atmospheric and ground spillover signals and thereby maximize their cancellation. We therefore adopt the same "double switching" technique introduced in Readhead et al. (1989) and also discussed in Myers et al. (1997), Leitch (1998) and Angelakis et al. (2009). In this method, in addition to the dual-beam Dicke switching, we also switch the target between the two beams and hence

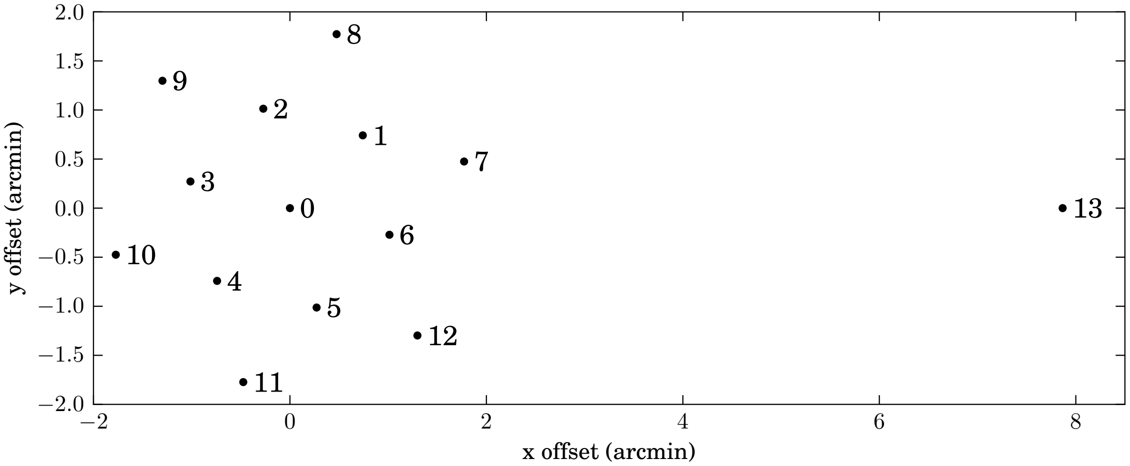

Figure 2.30. Pattern of offsets for a POINT2D procedure. The expected position of the source is at 0.

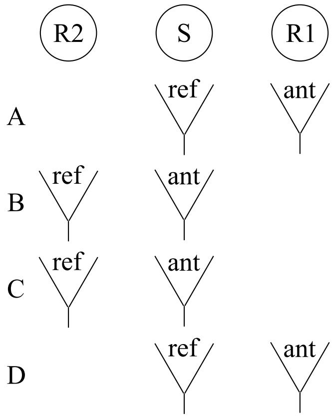

Figure 2.31. Schematic illustration of the FLUX procedure. The source (S) and reference fields (R1 and R2) alternate between the *ant* and *ref* in the four segments of the procedure.

remove both the constant term and any linear drifts in the power from these unwanted components of the signal.

FLUX procedure. The FLUX procedure implements double-switching and is the measurement procedure by which we collect the flux densities for this monitoring program. Figure 2.31 illustrates the procedure schematically. Each FLUX procedure encompasses four Dicke-switched integration periods of length τ, labeled A, B, C, and D. The FLUX measurement is executed as follows. First, the *ref* beam is positioned on the source and integrated, yielding the power difference $\xi_{\rm A}$. Then the *ant* beam is positioned on the source and integrated to produce $\xi_{\rm B}$. Next, $\xi_{\rm C}$ is measured with *ant* still on-source, and finally the *ref* beam is again positioned on the source for a final integration, $\xi_{\rm D}$. Thus we spend a total time of 4τ actually integrating on the source for each flux density measurement. Of course, slewing and settling times have to be allowed for at the beginning of the A, B and D integrations. In this program, $\tau = 8$ s is used and with slewing overhead, the total time for a FLUX procedure is about 1 min, yielding an on-source efficiency of about 50% for the FLUX procedure.

From the A, B, C, and D integrations, the corresponding flux density is given by

$$S_{15} = \frac{\kappa}{4}(\xi_{\mathrm{B}} + \xi_{\mathrm{C}} - \xi_{\mathrm{A}} - \xi_{\mathrm{D}}), \tag{2.43}$$

where κ is the calibration factor required to turn DUinto Jy, and the rms error is given by

$$\sigma_{15} = \frac{\kappa}{4}\sqrt{{\sigma_A}^2 + {\sigma_B}^2 + {\sigma_C}^2 + {\sigma_D}^2}. \tag{2.44}$$

The calibration factor consists of a relative calibration factor that is computed for each measurement (section 3.2.2.1) and an flux density calibration factor (section 3.2.2.3).

The four integrations also contain interesting information about the stability of the instrument and, more importantly, the atmosphere, during the observations. For each flux density measurement, we therefore also compute two other quantities—one that we call the "switched power," ψ, given by

$$\psi = \frac{\kappa}{4}\left(\xi_{\mathrm{B}} + \xi_{\mathrm{D}} - \xi_{\mathrm{A}} - \xi_{\mathrm{C}}\right), \tag{2.45}$$

and the other that we call the "switched difference," μ, given by

$$\mu = \frac{\kappa}{4}\left(\xi_{\mathrm{C}} + \xi_{\mathrm{D}} - \xi_{\mathrm{A}} - \xi_{\mathrm{B}}\right). \tag{2.46}$$

Both ψ and μ should be zero in the absence of gain or atmospheric drifts so we use these as a way of estimating such variations in our error model (section 3.2.3) and to reject badly contaminated measurements (section 3.2.1.6). The uncertainties in ψ and μ are clearly given by equation (2.44).

A third combination of the four integrations cancels the contribution of the source field. We call this the "source-nulled flux," and it is defined as

$$S_{\mathrm{null}} = \frac{\kappa}{4}\left(\xi_{\mathrm{A}} + \xi_{\mathrm{B}} + \xi_{\mathrm{C}} + \xi_{\mathrm{D}}\right). \tag{2.47}$$

In S_{null}, the source field is canceled, leaving behind a signal proportional to the difference in brightness between the two reference fields. This property makes it a good detector for astronomical confusion, of which solar interference is an example. We do not use this combination during normal observing or data reduction, but we use it in section 4.4 to detect solar interference.

In practice, the 1 ms samples recorded by the data acquisition system are accumulated into 1 s averages for the *ant* and *ref* beams. The uncertainties ${\sigma_i}^2$ are measured from the sample variances in the 1 s averages rather than from the 1 ms samples directly.

Impact of contamination. Although we will show in section 4.5 that bright radio sources at high galactic latitudes are rare enough that very few source or reference fields will be affected by confusion, it is important to understand the effect of such contaminants. A contaminating source near enough the intended target to appear in the main on-source beam will introduce an artificially high flux density measurement. Fortunately, the average beam for the on-source field (figure 2.6) is nearly circularly symmetric, so if the 40 m pointing is repeatable, such contamination will occur with a constant antenna gain in every observation regardless of the parallactic angle. As a result, although the flux density measurements may be spuriously offset, no or very little artificial variability will be inferred.

Confusion from sources found in the off-source reference beams will reduce the measured flux density for a source. Fortunately, because each reference field is only integrated half as long as the source field, the amplitude of the contamination will be reduced by a factor of two. However, contamination in the reference fields is more likely to result in false variability because the beams are no longer symmetric under parallactic angle rotations. We discuss the effects of confusion further in section 4.5.

2.2.2.3 Calibration Diode Measurements

The output of the CAL and NOISE diodes is measured via the CAL procedure, which is very similar in operation to the FLUX procedure. If a source is being tracked, the telescope first slews to an offset position so the antenna beams are (most likely) positioned on blank sky near the target. A four-integration A, B, C, D procedure is executed, but no slew is required. Instead, during the B and C integrations, the NOISE or CAL diode is enabled. The output of the CAL procedure is then given by

$$S_{\mathrm{CAL}} = \frac{1}{2}\left(\xi_{\mathrm{B}} + \xi_{\mathrm{C}} - \xi_{\mathrm{A}} - \xi_{\mathrm{D}}\right), \tag{2.48}$$

and the rms error by

$$\sigma_{\mathrm{CAL}} = \frac{1}{2}\sqrt{\sigma_A^2 + \sigma_B^2 + \sigma_C^2 + \sigma_D^2}, \tag{2.49}$$

with similar expressions for S_{NOISE} and σ_{NOISE}. Note that the normalization factor is $1/2$ rather than $1/4$ as for the FLUX—this is because the diode is only active for two of the four integrations of the CAL procedure, whereas one beam is on-source for all four integrations in the FLUXPROC procedure.

Because the offset position at which the CAL measurement is executed is near the position of the source being tracked, the background level of the radiometer due to receiver noise, CMB, atmospheric, and ground contamination should be similar to that of a subsequent FLUX procedure. This is convenient because it leads to a similar level of gain compression for a CAL diode measurement as for a FLUX measurement of the source. Thus, unless the source is extraordinarily bright, calibrating a measured source flux density by comparing it to a nearby CAL will implicitly correct for nonlinearity. This is not true for the NOISE diode—the NOISE diode produces so large a signal that it induces a significant change in compression. Neither the CAL diode nor typical program sources are bright enough to induce such changes.

2.2.2.4 Other Observation Procedures

Two other procedures are used for calibration purposes. The ZERO procedure is used to measure the output of the detector when the radiometer is disconnected. This is achieved via an RF switch that can connect the detector diode input to a matched termination instead of the radiometer output. This procedure is used to distinguish the amplitude of the detected signal due to power incident on the radiometer from the offset due to the back-end itself. Such an offset measurement is not required for the differenced output, ξ, since any such offset would be canceled by the subtraction.

The AVERAGE procedure was used in the VAX control system to integrate the switched, ξ, and total power, P, signals from the radiometer without performing sky switching through the FLUX procedure. Because the MCS control system constantly stores the raw output samples, an explicit procedure for this is no longer required.