

A Measurement of the Sunyaev-Zeldovich Effect in the Coma Cluster of Galaxies

Thesis by

Thomas Herbig

In Partial Fulfillment of the Requirements
for the Degree of Doctor of Philosophy

California Institute of Technology
Pasadena, California

1994

(Submitted September 20, 1993)

*To my parents,
who brought me to the New World.*

ACKNOWLEDGMENTS

I would like to thank my adviser Tony Readhead, who was of essential influence to my work and to my thinking, and whose profound humanity and good company I will miss very much. I also want to thank Charlie Lawrence, who was not only an Ersatz advisor when Tony could not be found, but with whom one can always have an argument about Bruckner. Steve Myers is the one really responsible for my interest in the microwave background, and thus shares blame for this thesis. He was also the one who would always inject a healthy dose of realism into my doings; nonetheless we managed to stay on very good terms.

None of this thesis would have been possible without the dedication and the knowledge of the staff at OVRO, especially Harry Hardebeck, Mark Hodges, and Russ Keeney. They kept two temperamental radio telescopes with their receivers at bay, and always remained good sports even when everything broke. Not to be forgotten is Wally, whose enthusiasm remains an inspiration. Jack and Jeff kept up our spirit at various critical times in the morning. Thanks also go to NRAO, TRW, and JPL for their efforts in developing and building the HEMT devices and amplifiers without which this project would have taken at least five years longer. The efforts of Sam Gulkis at JPL were vital in starting the collaboration that allowed this instrument to be built.

I would also like to thank Mark Birkinshaw, who has been observing the SZ effect for a long time; his was an important influence in many aspects of this work. Thanks also go to the two undergraduates, Niel Brandt and Jonathan Baker, who worked on some aspects of this project during the summer of 1992.

This thesis benefitted tremendously from the environment and the support of the astronomy department at Caltech. I enjoyed its friendly and concentrated atmosphere throughout my stay, and I will miss the company of many. My three buddies, José, Nick, and Wenge, were always a good reason to return from the observatory. Alan Dressler, Rashid Sunyaev, and many others helped me by their warm encouragement.

ABSTRACT

This thesis presents a measurement of the Sunyaev-Zeldovich effect in the Coma cluster of galaxies. The observed effect is $\Delta T_{0,obs} = -302.1 \pm 48 \mu K$, and was measured at a frequency of 32GHz with the 5.5m telescope at the Owens Valley Radio Observatory. Interpreted in the context of existing X-ray models, I derive a peak central SZ effect of $\Delta T_{0,peak} = -495.3 \pm 78.7 \mu K$. With this interpretation, I also calculate a value of $H_0 = 74.2 \pm 27.9 \text{ km s}^{-1} \text{ Mpc}^{-1}$ for the Hubble constant. This is the first successful measurement of the SZ effect in a nearby cluster, an important advance because of the much better potential for its astronomical and cosmological interpretation.

TABLE OF CONTENTS

Acknowledgments	iii
Abstract	iv
Table of Contents	v
1. Introduction	1
2. Physics of the SZ Effect	3
2.1. The Sunyaev-Zeldovich Effect	3
2.2. Applications of the SZ Effect	6
2.3. History of SZ Observations	8
3. The Instrument	13
3.1. Description	14
3.1.1. Site	14
3.1.2. Telescope	16
3.1.3. Receiver	17
3.2. Characterization	26
3.2.1. Radiometry Procedures	27
3.2.2. Root-2-ology	29
3.2.3. Bandpass	31
3.2.4. 1/f Noise Behavior	34
3.2.5. Dicke Switch Isolation	35
3.2.6. 60Hz Pickup	36
3.2.7. The Beam	36
3.2.8. Pointing	38
3.2.9. Zero Levels	39
3.2.10. System Stability and the Length of Flux Measurements	41
3.2.11. Ground Pickup	42
3.2.12. Non-linearity	45
3.2.13. Cal Diodes	48
3.2.14. Outstanding Problems	51
3.3. Calibration	52
3.3.1. System Temperature Measurements	52
3.3.2. External Cold Load	55
3.3.3. Calibration	58
3.3.4. Efficiencies	60
4. Observations	63
4.1. Choice of the Cluster	63
4.1.1. Telescope Capabilities	64
4.1.2. Requirements	65
4.2. Observing Technique	67
4.3. Observing Positions	73
4.4. Data	74

5. Analysis	77
5.1. Overview	77
5.2. Data Editing	78
5.2.1. Cal Editing	79
5.2.2. Overlap	79
5.2.3. Bad Weather Removal	80
5.2.4. Pair Completion	82
5.2.5. Editing Methods Not Used	83
5.2.6. Rejection Statistics	84
5.2.7. Robustness	85
5.3. Results	87
5.3.1. Correction Factors	87
5.3.2. Results	88
5.4. Tests of Data Quality	89
5.4.1. Averaging Procedures	89
5.4.2. Subsets	90
5.4.3. Ground Subtraction	91
5.4.4. Statistical Distribution	91
5.4.5. Variability with Parallax Angle	93
6. Discrete Radio Sources	96
6.1. Existing Surveys	97
6.2. Observations at 18.5 GHz	98
6.3. The Halo Source	100
6.4. Corrected Results and Undetected Source Allowance	100
7. Interpretation	102
7.1. The Coma cluster in X-rays	102
7.1.1. Description	102
7.1.2. X-ray Models	104
7.2. Derivation of the Hubble Constant	106
7.2.1. Theory	106
7.2.2. X-ray Model Predictions	107
7.2.3. Comparison with Observations—the Hubble Constant	110
8. Discussion	112
8.1. Systematic Astronomical Effects	112
8.2. Peculiar Velocity	113
8.3. Error Budget	114
9. Conclusions	115
References	117
Appendix A	120
Appendix B	124
Appendix C	133
Appendix D	138
Appendix E	140
Appendix F	142

1. INTRODUCTION

Before the recent detection of intrinsic anisotropy in the cosmic microwave background (Smoot *et al.* 1992) at a level of tens of microkelvin, two sources of anisotropy had already been known and observed. The dipole of the microwave background, at a level of several millikelvin, is merely a consequence of our motion with respect to the Hubble flow. At a level of several hundred microkelvin, the other is caused by the distorting influence of hot internal gas contained in massive clusters of galaxies. This source of anisotropy is called the Sunyaev-Zeldovich effect and has a number of important astrophysical consequences.

This thesis describes a measurement of the Sunyaev-Zeldovich effect in the Coma cluster of galaxies with the 5.5m telescope at the Owens Valley Radio Observatory. It is the first detection of the effect in a nearby cluster, which is of considerable advantage for the astrophysical interpretation of the result. It is also the first scientific result from this new telescope.

Measurements of the Sunyaev-Zeldovich effect have many applications, in particular as a probe of the hot gas atmosphere of a cluster of galaxies, but also as an independent means of measuring the absolute distance to the cluster. This latter measurement is possible through a comparison with X-ray data on the cluster atmosphere, and has been used in this thesis to determine the Hubble constant.

The Sunyaev-Zeldovich effect is a very small deviation in the brightness temperature of the microwave background radiation in the direction of a cluster of galaxies, and considerable care must be taken in the preparation of the instrumentation as well as during observations and analysis. For this reason, this thesis pays careful attention to the characterization of the instrumentation, in whose development I have invested much effort. However, I want to stress that this thesis does not emphasize instrumental development. Even during periods of intense instrumental involvement, the project always remained driven and motivated by its scientific objectives.

After an initial introduction to the Sunyaev-Zeldovich effect, its applications and the history of its observations in chapter 2, the thesis is roughly divided into two sections. Chapter 3 forms the first of those with a thorough description and characterization of the instrument.

This information is mainly intended as a reference to the 5.5m telescope, and can be skipped by most readers.

The second section begins with a description of the choice of the target cluster and a description of the observations in chapter 4. The data from two seasons of successful observations are characterized and analyzed in chapter 5. This chapter also contains a discussion of the editing algorithms. The direct observational results are presented there, along with tests of their reliability. Chapter 6 deals with the corrections made for discrete radio sources, and chapter 7 makes the connection between the observed Sunyaev-Zeldovich results and the X-ray data available on the cluster. For this, X-ray models from the literature are discussed and model calculations are presented. This chapter concludes with a calculation of the Hubble constant. The final chapter outlines some of the pitfalls of the Hubble constant's derivation and summarizes the results of the thesis.

Several appendices are given. Two of them contain preprint copies of papers connected with this work. The first of these is a brief account of the observations and their results which has been submitted as a letter to *Nature*. The second is an account of a major instrumental redesign that is in press in the *Proceedings of the IEEE*.

2. PHYSICS OF THE SZ EFFECT

2.1. THE SUNYAEV-ZELDOVICH EFFECT

Predicted in 1968 (Zeldovich and Sunyaev 1969), the Sunyaev-Zeldovich (SZ) effect is the result of an inverse Compton cooling process in which the hot gas contained within a cluster of galaxies interacts with the low-energy photons of the microwave background that pass through the cluster, impeded only by a very small optical depth. This interaction alters the spectrum of the microwave background radiation passing through the cluster, producing a signature that can be detected today. We can measure this degree of Comptonization by comparing the microwave background radiation that passed through a cluster of galaxies with that from points adjacent to the cluster, where the microwave background radiation has remained unaltered.

In the case of a photon gas interacting with a non-relativistic electron gas satisfying $h\nu \ll m_e c^2$ and $kT_e \ll m_e c^2$, the appropriate Boltzmann equation for the time evolution of the photon phase space density is approximated by the Kompaneetz equation (Kompaneetz 1957, see Rybicki and Lightman 1979 for a good derivation)

$$\frac{\partial n}{\partial t} = \frac{\sigma_T n_e h}{m_e c} \frac{1}{v} \frac{\partial}{\partial v} v^4 \left(n + n^2 + \frac{kT_e}{h} \frac{\partial n}{\partial v} \right)$$

where the phase space density is $n = c^2 I_\nu / 2h\nu^3$, a Lorentz-invariant quantity. Using the expression for the Compton y-factor

$$y \equiv \int_0^\tau \frac{kT_e}{m_e c^2} d\tau_T = \int_{t_0}^t \frac{kT_e}{m_e c^2} n_e \sigma_T c dt$$

and the energy ratio

$$\tilde{x} = \frac{h\nu}{kT_e},$$

the Kompaneetz equation takes on a much simpler form:

$$\frac{\partial n}{\partial y} = \frac{1}{\tilde{x}^2} \frac{\partial}{\partial \tilde{x}} \tilde{x}^4 \left(n + n^2 + \frac{\partial n}{\partial \tilde{x}} \right).$$

In our case, we are interested in what happens to an extremely low-energy photon gas of black-body temperature T_r , passing through a very high-energy non-relativistic electron gas, so that $h\nu \ll kT_e$, and therefore $\tilde{x}^{-1} \gg 1$. This means that we can neglect the n and n^2 terms in the Kompaneetz equation, so that

$$\frac{\partial n}{\partial y} = \frac{1}{x^2} \frac{\partial}{\partial x} x^4 \frac{\partial n}{\partial x},$$

where

$$x = \frac{h\nu}{kT_r}$$

is the usual dimensionless frequency.

If the photon spectrum is that of a black body, the photon phase space density becomes

$$n = \frac{1}{e^x - 1}.$$

As a result, the abbreviated Kompaneetz equation above becomes

$$\frac{\partial n}{\partial y} = \frac{xe^x}{(e^x - 1)^2} \left[x \frac{e^x + 1}{e^x - 1} - 4 \right].$$

In the case of small y -factors or small optical depths, this can be simplified to a difference equation

$$\frac{\Delta n}{n} = \frac{\Delta I_\nu}{I_\nu} = \frac{xe^x}{(e^x - 1)^2} y \left[x \frac{e^x + 1}{e^x - 1} - 4 \right].$$

Since

$$\frac{\partial n}{\partial T_r} = \frac{xe^x}{(e^x - 1)^2} \frac{1}{T_r},$$

the change in apparent black-body temperature becomes

$$\frac{\Delta T_r}{T_r} = y \left[x \frac{e^x + 1}{e^x - 1} - 4 \right].$$

In the Rayleigh-Jeans regime, where $x \rightarrow 0$,

$$\frac{\Delta T_r}{T_r} \cong -2y = -\frac{2\sigma_T}{m_e c^2} \int n_e kT_e dl.$$

This describes an apparent decrease in the brightness temperature of the cosmic microwave background due to this inverse Compton scattering process. The decrease is exactly proportional to the integral of the electron pressure along the line of sight from the microwave

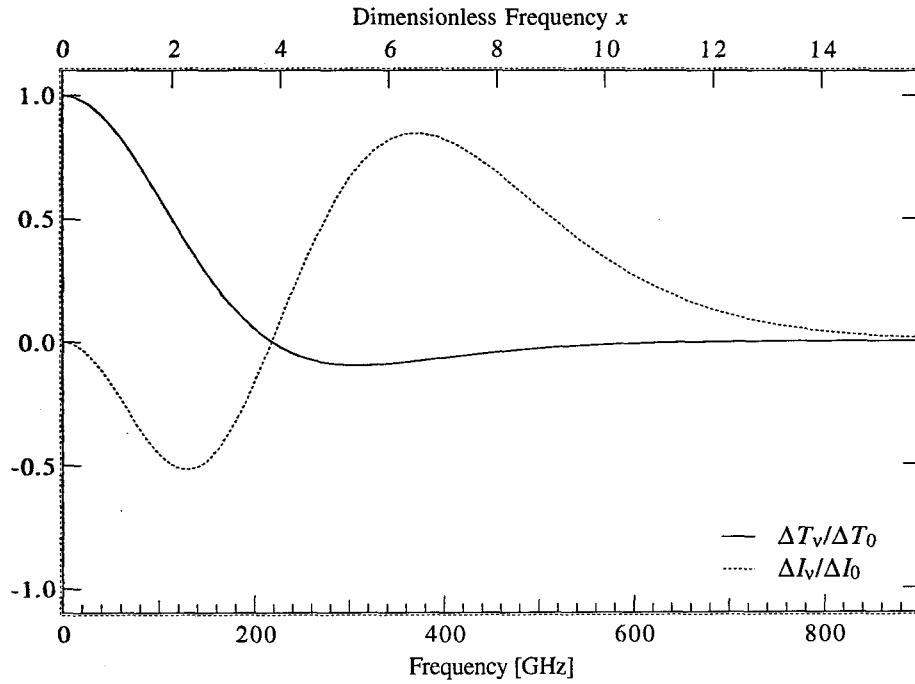


Figure 2.1 Frequency dependence of the SZ effect expressed as a temperature decrement and an intensity change. The normalization follows that of the text.

background's last scattering surface to the observer today. As can be seen from this integral, the size of the effect is independent of distance.

A look at the spectral shape of the SZ distortion is useful. This frequency dependence can be expressed in these compact forms:

$$\frac{\Delta T_v}{\Delta T_0} = -\hat{x}^2 \operatorname{cosech}^2 \hat{x} (\hat{x} \coth \hat{x} - 2)$$

$$\frac{\Delta I_v}{\Delta I_0} = \hat{x}^4 \operatorname{cosech}^2 \hat{x} (\hat{x} \coth \hat{x} - 2)$$

where $\hat{x} = x/2$. ΔT_0 corresponds to the zero-frequency strength of the effect, to which we will refer all observations later. ΔI_0 does not have such a physical interpretation, but it is useful as a normalization. Both relations are charted in figure 2.1.

The temperature decrement is strongest in the Rayleigh-Jeans section of the spectrum—it decreases until reaching the peak of the black-body curve at about 220GHz, where it becomes slightly positive. The intensity change, however, follows a very different path.

Zero in the Rayleigh-Jeans regime, it becomes negative at first with a peak at 170GHz, switches sign at the black-body peak and reaches a positive maximum at about 370GHz.

The SZ effect is small—even in a cluster with a large and hot X-ray atmosphere, the Rayleigh-Jeans brightness temperature decrement rarely reaches a peak of $\Delta T/T = 2 \cdot 10^{-4}$ in the line of sight through the cluster's center. At best, a realistic experiment will effectively measure one half of this amount, as seen in chapter 7. As a result, to achieve a high-quality measurement of the effect, experimental strategies very similar to those required for the study of intrinsic microwave background fluctuations must be employed.

2.2. APPLICATIONS OF THE SZ EFFECT

Perhaps the most incontrovertible application of the SZ effect is that its mere detection provides a fundamental proof that the cosmic microwave background originates from behind clusters of galaxies—this is an especially powerful argument with the SZ effect in distant clusters, such as the ones detected by Birkinshaw. Any argument invoking localized microwave background emission will inevitably run into trouble explaining the consistency of the SZ effects observed in clusters at redshifts as different as 0.02 and 0.55.

The Intracluster Medium

Quite independently from that, the effect can be used as a powerful probe of the intracluster medium of clusters at arbitrary redshifts. Since the SZ effect is merely a distortion in an extended screen of radiation, it is independent of distance, as long as it is resolved by the experimental configuration. The effect could be used to detect early, high-redshift clusters (*e.g.*, around high-redshift radio-quiet quasars), but also to find Zeldovich 'pancakes' in the early universe.

By the same token, since it is the strongest distortion of the microwave background radiation, it must be taken into account by studies of its intrinsic fluctuations. At the sensitivity levels now possible, it can be a serious contaminant for medium to high-resolution experiments (see Markevich *et al.* 1991, 1992; Bond and Myers 1993), but it could even play a role in low-resolution experiments such as COBE (*e.g.*, Hogan 1992).

More tangibly, SZ observations give a picture of the line-of-sight pressure integral through a cluster of galaxies. Such observations not only serve to constrain X-ray models of a cluster's X-ray atmosphere, but they also probe the cluster atmosphere to much larger radii than X-ray observations. The reason for this effect lies in the fact that the SZ effect depends linearly on the electron density, while the X-ray surface brightness, a two-particle interaction process, depends on the square of that density. As a result, the X-ray surface brightness profile of a cluster is much more strongly peaked than an equivalent SZ profile, while the latter extends to much larger radii from the cluster center.

Absolute Distances

A much more exciting application of the SZ effect lies in the reversal of this combination, allowing us to measure the absolute distance to a cluster directly. Such a measurement is independent of all other traditional distance measurements, relies solely on the calibrations of a radio and an X-ray telescope, and bypasses completely the well-known steps of the cosmic distance ladder. Combining this distance measurement with a nearby cluster's redshift, the Hubble constant is derived.

In principle, this distance measurement uses the fact that X-ray profile observations establish an angular profile of both electron temperature and density. Making the single assumption that the cluster's depth is equivalent to its width, this parameter profile can be used to infer an SZ effect from the cluster. Measurements of this SZ effect then provide the translation of angular size into physical length, and the cluster's absolute distance is derived.

Many derivations of this method have been given (see Sunyaev and Zeldovich 1980b), but the most general discussion comes from Birkinshaw (1991). It is outlined in section 7.2 and appendix F for the derivation of the Hubble constant.

The Velocity Effect

Lastly, there is another important application of the SZ effect, even though its use is still reserved for the future. If the cluster has some velocity relative to the Hubble flow, it will see a dipole anisotropy in the microwave background radiation field, just as we do on Earth. However, the fraction of the radiation field that is Comptonized will reflect the aver-

age radiation field experienced by the cluster, which is not altered by its movement. On transfer to our observing frame this fraction will be Doppler shifted, which can be observed as the peculiar radial velocity portion of the SZ effect:

$$\left. \frac{\Delta T_r}{T_r} \right|_{pec} = -\frac{v_p}{c} \tau_T,$$

where the positive sign of the peculiar velocity corresponds to recession (Sunyaev and Zeldovich 1980a). This effect is added directly to the main, or thermal, SZ effect. However, it can be distinguished by the fact that, unlike the thermal effect, it is independent of frequency. Therefore, high-fidelity observations of the SZ effect at several frequencies, spaced so that the thermal effect changes appreciably, can be used to measure the peculiar velocity of the cluster. It is also important to note that this application is completely independent of distance, and that it does not require any X-ray measurements. Its further virtue is that it is completely model-independent once the optical depth τ is known: given multi-frequency observations with exactly the same observing beam pattern, the peculiar velocity can in principle be obtained for any cluster, regardless of its symmetry or sphericity.

2.3. HISTORY OF SZ OBSERVATIONS

The formulation of the SZ effect prompted several programs to demonstrate its actual existence in clusters of galaxies. After producing a series of non-detections and unsubstantiated claims, they endowed the SZ effect with an aura of unreliability. This pessimistic perception is entirely unjustified today, considering the tremendous improvements in instrumentation and technique since those early attempts.

Perhaps most importantly, the sensitivity of the receivers of the time was strongly inadequate. For example, the high-frequency receivers of Lake and Partridge (1980), while advanced at the time, had a system temperature of 570K with a bandwidth of 1GHz. Comparing this to our systems, which combine a system temperature of 33K with a bandwidth of 5.7GHz, the sensitivity ratio is more than 30.

The consequences of this were manifold. Apart from reducing the necessary observing time, the most significant benefit of more sensitive systems is the improved ability to identify systematic instrumental effects. Ground pickup is one of the most pervasive of those systematic effects. Only with off-axis paraboloids or with pure horn antennas can this

pickup be reduced to an acceptably low level, but those antennas were and are still not easily available for SZ observations. Residual ground pickup effects can range from tens of mK to a hundred μK , too large to be ignored in the detection of the SZ effect (Parijski 1972; Perrenod and Lada 1979).

A further consequence of low-sensitivity systems is that observations tended to be made at lower frequencies, where better sensitivity was available. This exacerbates the contamination by discrete radio sources, most of which have steep power-law spectra. As Schallwisch and Wielebinski (1979) point out, coupling a low observing frequency with a high-gain antenna such as the Bonn 100m telescope, the signal is effectively masked in all but the most radio-clean clusters of galaxies.

More regrettably, the size of the expected effect was usually viewed with optimism. In the absence of high-quality X-ray data and well-developed observing techniques, detections were claimed at the level of several mK, only to be found spurious in later work (see Parijski 1972, Gull and Northover 1976, Lake and Partridge 1977, Perrenod and Lada 1979, and Birkinshaw *et al.* 1981). In addition, these expectations failed to take into account the reduction in observing efficiency that stems from the combination of the cluster's atmospheric profile with the effective beam shape of the experiment (see *e.g.*, Tarter 1978). This efficiency can be as low as 15–20%, but it rarely reaches above 60% of the central peak effect that can actually be observed by the telescope. Combining these factors, we find that even a cluster with a strong SZ effect like the Coma cluster can show less than a $-200\mu\text{K}$ deflection in the brightness temperature of the microwave background with an efficient observing configuration, far less than the optimistically claimed results from the 1970s.

Other experiments simply suffered from insufficient sensitivity to detect the effect (*e.g.*, Lake and Partridge 1980, Schallwisch and Wielebinski 1979, Rudnick 1978), and did not detect the effect in any of their clusters; in many cases their limits are consistent with modern data (Birkinshaw 1993 and this thesis), as well as with the experiment described here.

The first credible detection of the SZ effect was achieved in 1983 (Birkinshaw *et al.* 1984) with the OVRO 40m antenna and its K-band maser receiver, which was until recently the mainstay of the Caltech microwave background program. Among the several clusters observed at the time, three were found to show a strongly negative signal: Abell 665, the richest cluster in the Abell catalog, Abell 2218, and 0016+16. One of the main contributors to the solidity of those detections was the fact that the system's systematic effects and biases

had been carefully characterized in preparation for studies of the intrinsic fluctuations of the microwave background radiation at OVRO. Most likely, the level of systematic effects is less than $100\mu\text{K}$ in those data. Over the years, these detections could be repeated, along with good limits on a blank control field. In addition, observations were extended to form entire North-South profiles of measured points, which can be combined with X-ray profiles to form a self-consistent picture of the cluster atmosphere. Lastly, the observations were augmented by careful studies of the discrete source content of the clusters.

Confidence in microwave background results is strongly enhanced when they can be repeated with several telescopes. Using a similar configuration with a different telescope Uson (1986) confirmed Birkinshaw's detections in the three clusters. This encouraged many groups to pursue studies of the SZ effect with renewed vigor.

Recently, the SZ effect in A2218 was detected with the refurbished Ryle telescope at Cambridge (Jones *et al.* 1993), confirming Birkinshaw's detection in that cluster with a radically different experimental approach. First observations of the cluster 0016+16 with the OVRO mm-array at 32GHz are very encouraging as well (J. Carlstrom, private communication). Also, observations with the Australia telescope hold much promise for this approach (Ekers *et al.*, private communication). Interferometric observations of the SZ effect are a very interesting avenue, since they will eventually produce two-dimensional maps of the effect, and since they allow for very good removal of discrete radio sources. However, they are currently limited to observations of distant clusters because of the array antennas' primary beam sizes and their minimum separation, which resolve the effect in all nearby clusters. Because of the lack of small-baseline information in synthesis observations, single-dish and interferometric observations of the SZ effect will remain complementary.

Wilbanks *et al.* (1993), using a coupled bolometer array at the Caltech Submillimeter Observatory, recently detected the effect in A2163. This finding is especially interesting, since it is the first high-frequency detection of the SZ effect at 136GHz. Because of the spectra of most radio objects, such high observing frequencies bypass the effect of most discrete radio sources. Most importantly, however, they pave the way to observations of the peculiar velocity effect. This area of high-frequency observations will see much activity with balloon-borne measurements by Silverberg *et al.* (private communication) and future South Pole experiments by Dragovan *et al.* (private communication).

The radio source contents of A2163 has been studied in detail by Herbig and Birkinshaw at the VLA (1993, in preparation), who found a source with a strongly inverted spectrum in its very center. Predictably, subsequent observations with the OVRO 40m telescope showed no apparent SZ effect, despite the fact that A2163 is the highest-temperature cluster known with a temperature of $14\pm 1\text{keV}$ (Arnaud *et al.* 1992). The absence of a temperature decrement is entirely consistent with the influence of this contaminating radio source, but due to its strength and its most probably variable nature, its subtraction is not practical.

Because of their instrumental setups, most of these groups cannot observe the effect in clusters with the large angular sizes found in nearby clusters. The detected clusters are all at considerable distances, with redshifts from 0.17 to 0.55 for A2218 and 0016+16, respectively (Struble *et al.* 1987, Dressler and Gunn 1992). Only Silverberg *et al.* (1993) have attempted observations of the SZ effect in a nearby cluster with a balloon-borne bolometer, but failed to achieve the sensitivity required to detect the effect. Their target was the Coma cluster, and their upper limit is compatible with the detection described in this thesis (Silverberg, private communication).

Nearby clusters, ranging in redshift from 0.015 to 0.08 (see figure 4.1), hold particular advantages in the effect's interpretation (see chapter 7). Why are such nearby clusters not targeted by more groups, despite their advantages? One simple reason is that most high-quality steerable antennas are built with the aim of large surface area, a constraint that plays no intrinsic role in microwave background observations. Since most SZ observations were taken with existing telescopes and receivers, the differencing beam pattern on the sky does not extend far enough from the cluster center. Existing interferometers are affected by a similar constraint, which sets a minimum baseline length before shadowing or physical collision becomes a problem. Unfortunately, this makes interferometric observations even of small angular-size clusters such as A2218 a very model-dependent enterprise (see Jones *et al.* 1993).

Another reason for the difficulty of observing the SZ effect in large, nearby clusters are the problems inherent in large switching angles on the sky. In the case of mechanical switching, such as in the Silverberg experiment, the shortest switching time scale is limited by the mechanical constraints of a wobbling secondary reflector—rarely do those rates exceed 20Hz. However, even at a switching rate of 500Hz in our experiment, the telescope is sensitive to atmospheric instabilities much more than a similar system with a smaller switching angle. This was especially noticeable when comparing the data quality of the

OVRO 40m and 5.5m telescopes, which were observing at the same time during the 1993 season. During very good weather periods, both telescopes performed similarly well, but the data quality of the 5.5m telescope deteriorated much more rapidly than the 40m telescope when the weather worsened.

A role in the difficulty of observing nearby clusters was perhaps also played by the fact that there are no extremely rich clusters in our immediate vicinity (see Abell *et al.* 1989). Because richness is tied to the depth of a cluster's potential well, there is a tendency for richer clusters to possess larger and hotter X-ray atmospheres, which are in turn responsible for a strong SZ effect. Even the Coma cluster is much less rich than some of the very rich distant clusters. The nearest cluster of richness class 3 is A2670, with a redshift of 0.0745. On the other hand, the richest cluster in the Abell catalog, A665, has richness class 5 and lies at the much larger redshift of 0.182 (Sargent 1973). This cluster has indeed been observed to have a strong SZ effect (*e.g.*, Birkinshaw 1991). While this does not mean that nearby clusters cannot have strong SZ effects, they are generally weaker than those of richer clusters that can be found at larger distances.

3. THE INSTRUMENT

For the microwave background observations described here, the 5.5m telescope at the Owens Valley Radio Observatory was used. These observations constitute the first scientific result from this telescope.

Observations of the cosmic microwave background take place at the limit of our instrumental capabilities. Our aim to reduce systematic effects below $10\mu\text{K}$ using a telescope immersed in an environment of 300K requires an isolation of more than seven orders of magnitude. More than any other factor, the quality of the instrumentation determines the success of the project and the correctness of the result. For this reason, work on the instrumentation was an essential and major component of this thesis. This motivates a thorough description and characterization of the instrument in this chapter.

During the years of my involvement in the telescope's design and testing, numerous changes were made to the existing system. Some of these changes were of a major nature, such as the replacement of the original maser with a HEMT amplifier, the adoption of a direct-detection RF system, and the replacement of the telescope's secondary support leg structure to reduce ground pickup. In the description section below, I will describe the entire instrument in the state in which the observations were made. Some of the instrumental changes as well as tests of various components will be outlined and motivated in the subsequent characterization section.

Much of the discussion below is very detailed, especially in the characterization section. This level of detail is required to explain the suitability of the instrument for an experiment that is easily contaminated by severe systematic effects—this section can therefore serve as a reference to the telescope's performance.

3.1. DESCRIPTION

3.1.1. Site

The Owens Valley Radio Observatory (OVRO) lies at an altitude of 4000 feet in a desert valley that extends from north to south for about 100 miles. Bounded on the west by the massive range of the Sierra Nevada and on the east by the White and the Inyo mountains, it is isolated and shielded from major population centers and their radio emissions. Bishop, the nearest town, is the focus of a population concentration of about 10,000 people—certainly without threat to the observatory's radio environment. The only such threat comes from the U. S. Navy's China Lake test range to the south, from which we appear to receive occasional radio interference with the 40m telescope.

The observatory's location represents a very delicate and advantageous balance to microwave background observations. First, its altitude and its desert surroundings reduce the level of atmospheric water vapor significantly. Since water vapor is probably the most important influence on our observations apart from equipment and competence, such a low level is very desirable. However, there are some astronomical sites that possess even lower levels of atmospheric water vapor and overall less turbulent air. Examples of these are Mauna Kea and the South Pole. Both of those sites are also illustrations of the other side of the balance that is met very well at OVRO. Microwave background work must proceed at the 'cutting edge' of receiver and telescope technology. Not only does an established observatory provide a better environment for such work, but it is also easier to attract true experts to such surroundings. In addition, its relative proximity to Pasadena makes it very accessible to the microwave background research group there, a significant advantage given the importance of close interaction between the scientific and the engineering staffs on such a project.

There are currently three microwave background experiments at OVRO: the 40m telescope now spends most of the winter observing various aspects of the background, while the new 5.5m telescope is dedicated to these observations, possessing a system that was designed specifically for such work. The adaptation of the OVRO mm-array for microwave background work at 32GHz is nearing completion—it can provide two-dimensional microwave background data similar to such projects as the Ryle telescope, except that the

system will be more sensitive to large structures than existing high-resolution microwave background interferometers.

Model predictions of the atmospheric water vapor emission, calibrated in terms of antenna temperature, are shown in figure 3.1. This model is specific for the conditions and the altitude of OVRO and indicates the expected emissions during good weather, when the variable water vapor emission in the form of clouds is minimal. Integrated with the current bandpass of the 5.5m telescope, the effective atmospheric emission temperature for the 5.5m telescope is 10.1K at the zenith.




Figure 3.1 Predicted atmospheric emission temperature at OVRO (S. Gulkis, priv. comm.). The feature at 22.4GHz is produced by water vapor, while the steep rise toward higher frequencies is dominated by an oxygen line at 60GHz. Previously, microwave background observations took place to the left of the water vapor line—with the 5.5m telescope we are in the area between the two lines.

Such conditions are most frequently achieved during the winter; sub-freezing temperatures reduce the amount of emissive water vapor, and the stillness of the atmosphere provides an excellent basis for observation schemes that compare the emissions from closely adjacent patches of sky. Indeed, there are sometimes contiguous periods of weeks of such low-emission conditions, when a long-lived high-pressure ridge over the Great Basin redirects the winter storm track to the north or the south of the Owens Valley. During such periods, the temperature remains below freezing even during the day, even though the sun heats the ground unremittingly from a clear blue sky; the relative humidity close to the ground may stay below 10% during such a period.

Unfortunately, such periods are rare, and equipment problems are very disrupting during such times. None of the data described here were taken during pristine weather conditions—the end of California’s drought with an uninterrupted procession of winter storms

was not welcome from this perspective. Still, useful data could certainly be gathered, at night and during shorter periods of cold and clear weather conditions.

3.1.2. *Telescope*

Although the 5.5m telescope is a recent addition to the collection of antennas at the Owens Valley Radio Observatory, the physical telescope was constructed in 1968. It had served for many years as an interferometer component at JPL's Table Mountain observatory. As the result of our collaboration with JPL, the telescope was moved to OVRO for its present purpose.

The dish is a commercial type built by the Whittaker corporation with a diameter of 18 feet. The surface consists of 16 radial segments, leaving a central hole of 2 feet diameter. The front of these panels is covered by a solid aluminum skin that is attached to the panels' internal honeycomb structure with an adhesive. No rivet or screw heads disturb this surface. The panels abut each other tightly, leaving a gap of typically less than 1mm. The backup structure of the dish is very rigid. As a result, the antenna surface does not deform measurably as the telescope tilts, so that the antenna gain remains constant at all zenith angles.

A hole with a diameter of 2 feet cuts through one of the top panels at a radius of about 2 meters from the apex. It is used by a television camera whose optical axis is aligned approximately with the telescope's radio axis, and which serves to find the initial radio pointing offsets. During normal observing, this hole is covered by an aluminum plate to prevent ground spillover.

The telescope is mounted in a standard altitude-azimuth (*altaz*) configuration for our experiment. Since it was originally mounted equatorially, the geometry of the counterweight arms is peculiar and makes balancing the telescope difficult at some zenith angles.

The telescope pedestal is a commercial type as well, built by Scientific Atlanta in 1968, and uses two variable-speed DC motors for each axis. In slew mode, both motors have the same direction, while in track mode the directions are opposite to eliminate backlash and to allow slow and accurate motion of the dish. When moving to a position, the telescope first moves there at high speed in slew mode. Once the telescope has come within 2° of the target, the drive system switches to track mode and pulls slowly toward the requested position, reducing the track speed as it closes in. In this way, the telescope overshoots

only minimally during acquisition, saving valuable observing time. It is also helpful that the telescope slews very fast; it can complete a 360° turn in less than two minutes.

The angular encoders are thermally controlled Baldwin encoders of somewhat advanced age that use a Germanium photodiode array illuminated by a flash bulb twice a second. This determines the frequency of the telescope drive servo cycle. Because of the encoders' age and its Germanium electronics, encoder problems were encountered frequently during observations and seemed to have a predilection for good observing weather, especially during the 1991-92 season. The encoder resolution is $0.^\circ005$ (18 arc seconds).

The guiding design consideration for the telescope's optical system was the minimization of ground pickup by the antenna, within the constraints of an obstructed aperture. As had been the case for the solar interferometer, we use the dish in a Cassegrain configuration to accommodate the large and heavy receiver and to reduce direct ground spillover. The hyperbolic secondary has a diameter of 2 feet and is constructed of cast aluminum. During the replacement of the feed legs, this mirror was retrofitted to contain a vertex plate which dramatically reduces the illumination of the center of the dish, where the receiver front and its mounting brackets can scatter the radiation in unwanted directions.

The secondary support leg structure is a tripod whose legs are led over the edge of the dish. The configuration is an inverted Y. The angle separating the bottom legs is 90° , giving an angle of 135° to the top leg. The width of the legs in the radiation direction is 1.5in, giving very low geometrical obstruction. The dish edge of the legs is modified by a baffle arrangement to direct any secondary reflections away from the ground. This feed leg arrangement is described in detail in appendix B.

3.1.3. Receiver

Design

While the telescope portion of the instrument utilized and adapted an existing telescope, the receiver of the 5.5m telescope was designed specifically for observations of the cosmic microwave background. Originally built as a heterodyne receiver with a low-noise Ruby

maser as its first-stage amplifier, its heart now consists of a low-noise HEMT amplifier that far exceeds the sensitivity of the original maser.

In addition, the receiver now has a continuous RF signal path until detection, without any frequency downconversion. There were several reasons for this design, but perhaps the two most important ones are the smaller fractional bandwidth of a direct-detection system and the removal of the local oscillator system, which introduced considerable 60Hz interference into the system.

Other considerations affected the design of the receiver at a fundamental level. Particular attention was paid to the ease of calibration, but also to the elimination of potential systematic effects. For this reason, the electronic backend after detection is fully digital and samples at a frequency of 1kHz. No analog signals leave the shielded environment of the receiver's outer can. Many system parameters, such as temperatures and pressures, are monitored continuously and are available to the control and acquisition computer. A total of 32 such monitor channels are available.

Because of various constraints, unlike some other low-frequency microwave background systems (*e.g.*, Gaier 1992), our receiver consists of only one frequency channel. While this limitation is a consideration for intrinsic microwave background anisotropy work, it does not adversely affect observations of the SZ effect.

As can be seen in the block diagram of figure 3.2, the receiver is based on a standard Dicke switched radiometer. The Dicke switch combines two arms that are identical not only in their components, but also in their physical layout. After that, there is only a single signal path, all the way into the computer. No component in the entire receiver (except for the Dicke switch) is aware of the Dicke switch position—this is valid even for the data section of the backend. The Dicke switch position is only used by the computer and is applied to the incoming samples, with no influence on a sample's data value.

The entire receiver package is kept within a 10K temperature range by a heating and cooling air handler.

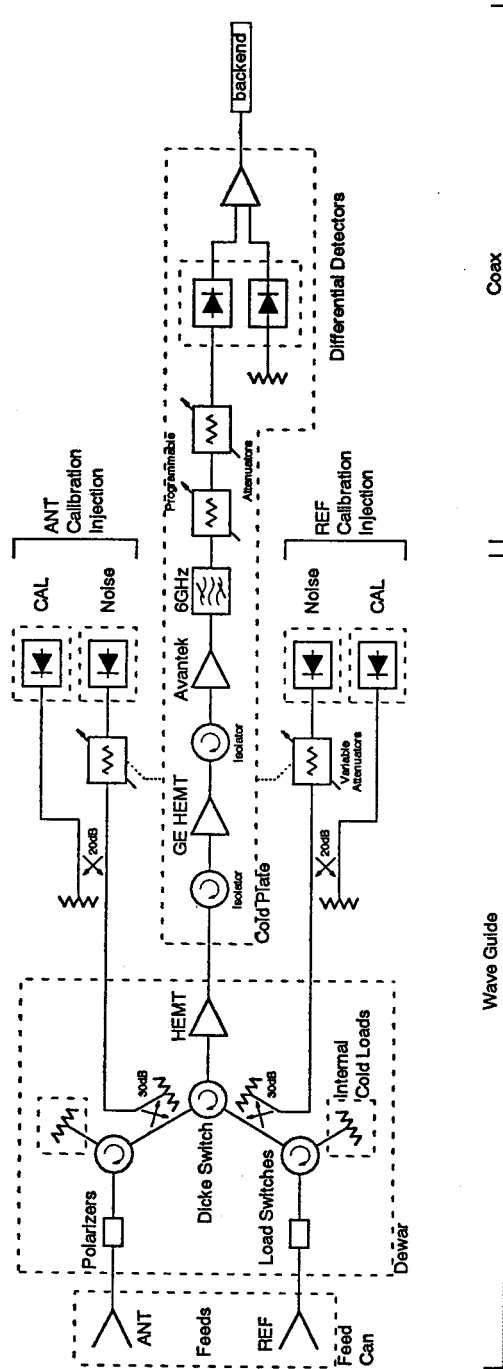


Figure 3.2 Schematic block diagram of the 5.5m telescope's receiver. See text for details

Feeds

The receiver's frontmost elements are the two identical circular corrugated feeds that illuminate the secondary mirror with such a taper that the power at the secondary's edge is reduced by 20dB. Because of the Cassegrain configuration, the feeds must illuminate a small angle and are hence quite large: their length is 18in. The front perimeters of the feeds touch each other; this sets the beam separation on the sky. Because of their size, the feeds are not part of the dewar. Since feeds inherently introduce about a 0.1dB loss, this results in an increase of the receiver noise temperature by 3–5K. To avoid condensation on the feeds, they are enclosed in an aluminum can with a Kapton front window; this can is pressurized at 5psi with nitrogen gas. Inside this forward can a small fan circulates air to equalize the two feeds' temperatures and hence their noise power contributions.

Cryogenic RF Section

The circular waveguides from the feeds enter the dewar through vacuum windows and an infrared radiation block, and are converted to standard WR28 waveguide after the left-circular wide-band polarizers. After this stage, each arm encounters a circulator switch (identical to a Dicke switch, but not switched continuously) that allows one to select either the incoming sky radiation or radiation emitted by two identical thermally isolated and heatable loads. The loads are typically at a physical temperature of 17K. They are not packaged together; this means that most of their small temperature variations are uncorrelated and ultimately limit the use of these loads for long-term stability measurements. Another consequence of these internal cold loads is that the loads radiate out of the feeds in the sky position—this is of concern for the calibration with external loads. Overall, however, these internal cold loads have been a vital tool for the telescope's testing.

After calibration noise injection via 20dB cross-guide couplers, both arms finally join at the Dicke switch. The Dicke switch and the Load switches are identical wide-band ferrite waveguide switches with an insertion loss of 0.2dB or better. In our system the bandwidth is limited by the 20% bandpass of these switches. These modern ferrite switches are much improved over earlier types; they are compact and because of their low power consumption can be switched rapidly without warming measurably. In addition, they can switch very

quickly: we have measured typical settling times of less than $40\mu\text{s}$ when switching between two very different signal levels. The combination of these properties allows us to switch the Dicke switch as often as every millisecond.

The dewar is cooled by a two-stage Gifford-MacMahon CTI 1020 refrigerator supplied with 300psi helium from a compressor stationed on the telescope's concrete pad. The 70K stage and its shield is only used as an intermediate station to separate the cold stage from the outside. All RF components are mounted on the 15K stage, thermally isolated from the outside via copper flashed stainless steel waveguide sections. The ease of operation of a 15K system is made possible by the cryogenic performance of HEMT devices whose gain and system noise temperature are virtually independent of temperature below about 25K. Most certainly, our observing efficiency has improved markedly over that achieved with masers because we could discard the finicky 4K system required by those amplifiers.

HEMT

The signal then enters the first-stage amplifier, the receiver's heart. This receiver was originally built as a prototype for a new series of K_a band masers at JPL. This maser, however, never lived up to its full potential—on the telescope, the full bandpass of 400MHz and the desired receiver noise temperature of 15K were never achieved. This was partly due to an insufficient cryogenic system, which could not handle the enormous load of the maser and could **only** keep the system at superconductive temperatures when pointed at the horizon. As soon as the telescope was raised from that position, the system would warm up and the maser would lose its magnetic charge. In addition, the Impatt maser pump diodes were very temperature sensitive and injected an enormous amount of RF power into the system. This situation was intolerable.

Since September 1990, we have used High Electron Mobility Transistor (HEMT) amplifiers for this first-stage amplification; because of their bandwidth, their low noise, their stability, compactness and cryogenic tolerance, these amplifiers are now the standard for all advanced broad-band cm-wave receivers. Installation of the HEMT and removal of the maser simplified the system enormously and reduced the cryogenic load considerably.

Our first HEMT was purchased from the National Radio Astronomy Observatory's (NRAO) advanced receiver lab in Charlottesville and could be used in our system without modification. Its four stages, each consisting of a Fujitsu GaAs transistor and some supply

and matching circuitry, give a gain of 27dB with an amplifier noise temperature of about 45K over the band we use. Even judging by the specifications only, this HEMT gives better performance than the promised maser system: a sensitivity of $0.6 \text{ mK s}^{-1/2}$ as compared to the maser's $0.75 \text{ mK s}^{-1/2}$. The real improvement, however, is much more considerable. When comparing the system to the successfully working K-band maser system on the OVRO 40m telescope, the increase in reliability is significant. While the cryogenic system with a maser may achieve continuous run times of no more than a month before a new cool-down and its down time is required, we customarily run the HEMT cryo system for the entire observing season without observing any signs of deterioration. In addition, a maser's gain is exponentially dependent on the physical amplifier temperature, which has cost much observing time due to cryogenic temperature oscillations in the past. In terms of overall observing efficiency, I would judge this HEMT system to be more effective than a working maser system by a factor of at least two to three.

We have also tried to keep pace with the rapidly progressing HEMT technology. In the summer of 1992 we upgraded our amplifier to a system that was assembled at JPL's receiver lab, containing as its first stage an experimental indium phosphide device developed by TRW and provided by the courtesy of Barry Allen. This new amplifier has a noise temperature of about 20K over the entire band; this translates to a sensitivity of $0.26 \text{ mK s}^{-1/2}$, three times the expected sensitivity of the maser system.

Upon arrival of this new amplifier at OVRO, we minimized the system noise by adjusting the amplifier's bias settings. While the bandpass and the overall gain is not strongly variable within a narrow window of bias parameters, the noise performance is. This is on contrast to my experience with the NRAO-built HEMT amplifier that replaced the 40m telescope's K-band maser: there, the noise temperature hardly varied with bias parameters, but the gain and the bandpass could be affected noticeably.

Cold Plate

Because of the high gain of the first-stage amplifier (about 27 or 32dB for the NRAO and TRW amplifiers, respectively), all further components can be operated at room temperature outside the dewar. Still, all components outside the dewar and before the backend's integrator are mounted on a thermally insulated 1/4in copper plate that is controlled to 0.01K via an electrothermal heater/cooler element. As can be seen in figure 3.2, the

remaining RF chain consists of two more amplifiers for 55dB more gain, separated by isolators, a 6GHz bandpass filter matched to the bandpass of the Dicke switches (29–35GHz), and two programmable HP attenuators (one with 1dB steps for 11dB total, the other with 10dB steps for 90dB total). The system shows some non-linearity before the attenuators; we currently suspect the second room-temperature amplifier to be its source. This will be discussed in detail below.

The detection of a 6GHz bandpass at 32GHz is not entirely straightforward, since the sensitivity of Schottky barrier diode detectors is very low at such high frequencies. To boost this sensitivity, the diodes are forward-biased at 100 μ A. Unfortunately, the output offset of biased detectors is much more temperature sensitive than we can afford, so we combine two identical detectors to form a differential detector pair: one detector connected to the signal path, and the other's input terminated by a load (which is very close to zero signal). Both detectors are biased identically and are in intimate thermal contact with each other, with their output differenced by high slew-rate operational amplifiers. This scheme reduces the thermal detector offset drift dramatically, but does not eliminate it entirely, as discussed below. Another consequence of this use of a biased detector is the introduction of some additional nonlinearity, which is discussed below as well.

Backend

The output of the detectors is sent to the backend via a differential line to cancel common modes. The backend resides within the receiver package and is fully digital; its only analog components are the integrator and the analog/digital (A/D) converter. The integrator-A/D circuitry consists of precision commercial components and is highly linear. Considerable care was taken to separate it from the backend's digital circuitry. The original Burr-Browne A/D is a true 16-bit converter, but was replaced in the summer of 1992 after it failed. The replacement converter, another true 16-bit circuit from Analog Devices, has better linearity and digital ground rejection, and sports an explicit self-calibration feature that was invoked twice a day during our observations in the 92–93 season. This self-calibration allows for additional precision in the A/D converter's linearity.

The fundamental sampling cycle is 1ms and is built into the backend. Combination into longer integration times is handled by the data acquisition and control computer. Between sampling integrations, the integrator is blanked for at least 10 μ s to allow the system to

settle. However, if the Dicke switch changed position from one integration to the next (the standard case for our observations), an additional blanking period is applied. We found that a total blanking of $48\mu\text{s}$ provided a good margin for the rejection of Dicke switch transients. Note that the 16-bit A/D converter always measures the total radiometry power for the current sample—it does not measure switched power directly.

The digital backend also handles all 32 incoming monitor channels from the receiver: one of them is converted by a separate 12-bit A/D every millisecond. The monitor points are multiplexed, so that each point will be sampled every 32ms.

Finally, the backend receives hardware settings instructions from the control computer; all of them can be changed every millisecond. These settings, such as the Dicke switch position, are buffered by the backend and are returned to the control and acquisition computer along with the radiometry and monitor samples during which they were active. With this scheme, the computer has full settings information for each 1ms sample, making the backend robust against timing glitches.

Calibration Capabilities

Calibration and stability are two of the most critical components of any microwave background telescope. For this reason, the 5.5m telescope's receiver provides ample calibration capabilities. First, it possesses internal cold loads on both arms as described above. These loads can be heated slowly, over the course of 20 minutes, their temperature being monitored by precision Germanium resistors, to provide a variable power level. However, these serve mainly to provide a stable, low power source as a background for other calibration procedures.

Each arm has two calibration noise injection capabilities: a large one to provide a signal comparable to the system's noise temperature; the other small to provide a signal comparable to sources in the sky. The large noise sources can be attenuated with a variable Hughes waveguide modulator, where both arms are controlled by the same computer setting simultaneously.

The small noise sources, called the CALs, are the workhorses of the calibration. Since they are not variable and since they are small enough to be used during regular observations, all calibration was referenced to them. The four noise sources are diodes that are supplied by four independent constant-current sources in the backend that are on constantly. The

diodes are then turned on and off by a sink transistor that blocks or sinks the current source to ground, respectively. Thus the noise diodes can be turned on and off very rapidly, without having to wait for the current sources to stabilize.

Control and Data Acquisition

The telescope is controlled by a disk-less microVAX computer that acts as a single-user satellite to the standard OVRO computer control program (CCP). The satellite also handles all data acquisition. It is here that all radiometry results from the backend are averaged in bins corresponding to the Dicke switch position; this also makes the system's output robust against timing glitches. The control and acquisition software is described in detail in the *OVRO 40m and 5.5m Control System* manual (Pearson 1993).

Apart from controlling receiver settings, the control software initiates procedures that can accomplish tasks from simple integrations to complex skydips; the most commonly used radiometry procedures are described in detail below. All data from the telescope (which are results of procedures) are recorded in a text-format log file. Since the data rate for single-dish microwave background observations is relatively low, no special data files are created. Analysis programs must be able to read the complex syntax of these log files. A significant advantage of this data collection is that all data appear in their observing schedule's context, since all scheduler commands are echoed in the same file, along with status and error messages, as well as real-time user comments about weather and instrument conditions. This greatly facilitates the subsequent data analysis.

When tracking a source in the sky, the telescope constantly updates its position to keep the sky source in the center of the ANT beam.

Chart Recorder

The control system provides for 8 channels of chart recorder data: these channels can be any combination of the 32 monitor channels and the 4 radiometry channels, and may be selected through the control software. Two of these channels can be charted by a conventional chart recorder, but the full 8-channel capability is provided by a multi-channel Astromed chart recorder. Located in the control building, it is usually run in 8-channel mode, recording 4 channels from the 5.5m telescope and 4 equivalent microwave back-

ground channels from the 40m telescope. This affords a real-time visual comparison of the two systems and their weather sensitivities. The Astromed chart recorder is initialized by the control computer and produces continuous sheets with a thermal printer; both the traces as well as the grids are printed by the recorder itself. In addition, each trace is labeled and the time and date is recorded near the paper's edge. Remembering the old chart recorder rolls whose timing had to be marked by hand (which was all too often forgotten), this recorder system truly is a leap forward.

Alarms

Recently, alarm capabilities were added to the computer control system. Alarms include motion alarms, receiver communications alarms, as well as user-settable alarm ranges on any of the 32 monitor channels and the 4 radiometry channels. Not only does an alarm produce a record in the telescope log file, but it can trigger a physical alarm as well. I have provided an alarm controller that can relay such an alarm to a paging company, ringing a pager as far away as Pasadena to alert the observer.

3.2. CHARACTERIZATION

Over the years, extensive effort went into thorough testing of the instrument. These tests establish the viability of the 5.5m telescope for microwave background observations. For this reason, and because these tests have occupied a significant part of my thesis work, I have chosen to explain the tests and their results in detail below.

The SZ observations described here were made with two different RF amplifiers. While they are not fundamentally different, many system parameters have changed. In particular, the bandpass, the zero levels, the non-linearity, and the calibration changed from the 1992 season to the 1993 season. This will be noted in the appropriate sections.

3.2.1. Radiometry Procedures

Before delving into the details of the instrumental tests, it may be useful to review the most common types of data acquisition procedures we use. In all of our observing procedures involving switched radiometry, we used a full Dicke switch period of 2ms.

In the description of the backend I mentioned that the fundamental sample has a length of 1ms. These individual samples are collected by the acquisition computer and are tallied by Dicke switch position. These tallies form secondary samples that are then used for longer integrations in the standard radiometry procedures. The length of secondary sample intervals is user-selectable; we have kept it at 1000ms for observing purposes, and at different (usually shorter) lengths for some test procedures. The standard radiometry procedures do produce an estimate of the internal statistical reliability of their results. This estimate is the result of forming the standard error of the accumulated secondary samples, which themselves do not retain any statistical information from the incoming 1ms samples. While that information would be desirable to improve the quality of the uncertainty estimate and to reject individual secondary samples (*e.g.*, containing a 20ms burst of interference), it is not computationally feasible considering the capacity of the acquisition computer and the demands placed on it by telescope control. The internal error estimate thus contains information on the statistical fluctuations on time scales of about 10 seconds.

Because of the hardware integration scheme, the telescope's backend provides total power readings only. Switched measurements are formed by subtracting the average of the total power samples with the Dicke switch in the REF position from those with the switch in the ANT position. This means that four radiometry channels are available for recording: the total ANT and REF powers, the switched power $P_{sw} = P_{Ant} - P_{Ref}$, and the average total power $P_T = (P_{Ant} + P_{Ref})/2$.

Average

This procedure is the simplest one of the radiometry procedures. It merely averages S secondary samples in 1–8 simultaneous channels, which can contain any of the 4 radiometry channels as well as any of the 32 system monitor channels. The procedure will idle for I secondary samples before starting the integration. The timing for the procedure is $\tau_A = I\tau + RS\tau$, including an allowance for several repeats R .

Zero

The Zero-level procedure also performs a simple integration, but it integrates the two total power radiometry channels at full RF attenuation. This measures the system's zero level offset, which originates from any of the DC elements from the detectors on backwards. This zero level is of interest in all total power measurements. Similar to the average procedure, the timing for the zero-level procedure is $\tau_z = I\tau + RS\tau$.

Flux

The Flux procedure is the workhorse of the observations. It performs a standard double-switched flux measurement (see *e.g.*, Readhead *et al.* 1988, and below) by measuring the Dicke-switched value in four segments, A through D. During segments B and C, the nominal position (usually containing the source of interest) lies in the center of the ANT beam, while it is placed in the center of the REF beam during segments A and D. The telescope has to move between segments A–B and C–D, which does take some time. This period may be of varying lengths (*e.g.*, longer in the A–B direction than in the reverse, C–D, direction). To avoid possible systematic effects from such an asymmetry, there is an *idle* parameter I , specifying the number of secondary samples to wait before attempting to start a segment integration after one of these moves. If the telescope is in track mode (the standard case during observations), a segment integration will not start unless the center of the ANT beam is within 30" (user-selectable) of the nominal sky position, including any applicable offsets. There is no time gap between segments B and C.

Including the repeat parameter R for multiple consecutive executions of the Flux procedure and the segment integration time S , the timing becomes $\tau_F = I\tau + R(4S\tau + 2I\tau)$.

Cal

Reference to the value of the calibration noise source is achieved with this procedure. Similar to the Flux procedure, it also consists of 4 segments, but no telescope movement is involved. Segments A and D are normal switched integrations, while segments B and C are identical integrations with one of the CAL noise sources switched on. The side of the noise injection can be chosen, but for all observations we have used the CAL on the ANT side.

A particularly useful feature is that we have the choice to perform a Cal procedure on the internal cold loads. This means that the Cal result is isolated from observing weather conditions to first order.

The result of the Cal procedure is a value for the injected noise power in instrumental counts. Since our external calibration establishes a (constant) value for the power of the noise source in antenna temperature units, we can use the Cal procedure to obtain the current calibration of the instrument in units of antenna temperature.

Just like in the case of the Flux procedure, the Cal procedure's timing is $\tau_c = I\tau + R(4S\tau + 2I\tau)$. A practical difference is the much shorter idle parameter, since the telescope does not have to perform any positional moves, and since the injected noise level settles in less than one second.

3.2.2. *Root-2-ology*

The use of switching radiometers often gives rise to confusion over the various factors of 2 that should appear in the standard radiometry equation. Below, I will give an exhaustive if tedious explanation of the sources of these factors. In principle, every level of subtraction produces a factor of 2 in the radiometry equation: one factor of $\sqrt{2}$ comes from the fact that one spends only half one's time observing the primary quantity; the other factor of $\sqrt{2}$ arises from error propagation through the subtraction operation. The latter factor is premised on the similarity of the noises on both sides of the subtraction; this is not an unreasonable assumption since most subtraction schemes strive to achieve this balance. Here are the details.

Single Horn Receiver

For a receiver with one channel, system temperature T_{sys} , and bandwidth $\Delta\nu$, the uncertainty in the brightness temperature estimate from an integration of time τ is limited by thermal noise in the usual radiometry equation

$$\sigma_T = \frac{T_{sys}}{\sqrt{\tau \Delta\nu}}.$$

Here, T_{sys} is the total system temperature including instrument, ground pickup, atmosphere, and sky emission, such as the cosmic microwave background radiation and radio sources.

Dicke Switched Dual Horn Receiver

For a receiver that Dicke switches between two channels ANT and REF, the differential brightness temperature between the two arms is

$$T_S = T_{Ant} - T_{Ref}.$$

From this follows that the thermal noise in the measurement is

$$\sigma_{T_S} = \sqrt{\sigma_{T_{Ant}}^2 + \sigma_{T_{Ref}}^2},$$

where the σ_{T_i} are the uncertainties in the estimates of the brightness temperature in the respective channel. Usually, $\sigma_{T_{Ant}} \approx \sigma_{T_{Ref}}$, so that

$$\sigma_{T_S} = \frac{\sqrt{2} T_{sys}}{\sqrt{\tau_{Ant} \Delta\nu}}.$$

Since $\tau_S = \tau_{Ant} + \tau_{Ref} = 2\tau_{Ant}$,

$$\sigma_{T_S} = \frac{2 T_{sys}}{\sqrt{\tau_S \Delta\nu}}.$$

Double Switched Fluxes

In a standard Flux procedure, four integrations A, B, C, and D are performed, each of which are Dicke-switched integrations. The sky temperature from this procedure is calculated as

$$T_F = \frac{1}{2}(T_B + T_C) - \frac{1}{2}(T_A + T_D).$$

It follows that the thermal noise in the flux measurement is

$$\sigma_{T_F} = \frac{1}{2} \sqrt{\sigma_{T_B}^2 + \sigma_{T_C}^2 + \sigma_{T_A}^2 + \sigma_{T_D}^2}.$$

Since again $\sigma_{T_A} \approx K \approx \sigma_{T_D}$, this becomes

$$\sigma_{T_F} = \sigma_{T_A} = \frac{2T_{sys}}{\sqrt{\tau_A \Delta\nu}},$$

where τ_A is the integration time of one segment of the procedure. Because the total integration time for the procedure is $\tau_F = 4\tau_A$,

$$\sigma_{T_F} = \frac{4T_{sys}}{\sqrt{\tau_F \Delta\nu}}.$$

This can be used conveniently to calculate the system temperature

$$T_{sys} = \frac{1}{4} \sigma_{T_F} \sqrt{\tau_F \Delta\nu}.$$

The system temperature calculated from this is usually larger than the nominal value, and is called the *dynamic system temperature*. Because of the fluctuation spectrum of system noise, its size is a function of the integration time scale, and is explored below.

If we arrange for the Flux procedure to have the sky source in the ANT channel for integrations B and C and to have the source in the REF channel for integrations A and D, then the source antenna temperature is

$$T_{source} = \frac{1}{2} (T_F \pm \sigma_{T_F}).$$

3.2.3. Bandpass

The receiver's bandpass has been one of the foci of our efforts to increase the system's efficiency. This has taken two forms: first, we tried to achieve the maximum nominal bandwidth permitted by the front elements, notably the Dicke switch; second, we attempted to smooth the bandpass to approximate the maximum possible effective bandwidth (*e.g.*, Herbig and Readhead 1991).

In sensitivity calculations, it is customarily assumed that the bandpass is rectangular and that it has a width of $\Delta\nu$. If the bandpass is not flat, however, the samples from some sections of the spectrum are weighted more than those from others. This will reduce the sensitivity and diminish the effective bandwidth of the system.

In the general case, the effect of the bandpass shape on the sensitivity is given by D , the degradation of the signal-to-noise ratio compared to the case of a nominal, flat bandpass of width $\Delta\nu_{nom}$ where (*e.g.*, Thompson and D'Addario 1982)

$$D = \frac{\int_0^{\infty} P_v d\nu}{\left(\Delta\nu_{nom} \int_0^{\infty} P_v^2 d\nu\right)^{1/2}},$$

and where P_v is the specific power (per unit frequency) in the bandpass. Thompson and D'Addario observe that a loss of only 2.5% in sensitivity occurs if the bandpass has ripples of amplitude 2.9dB, or if it has an overall slope of 3.5dB.

The bandpass was measured for both amplifiers and then integrated to determine the effective bandpass. In these measurements, a test signal was injected through the ANT calibration cross-guide coupler while the receiver was in the laboratory. Unfortunately, this arrangement bypasses the feed, the polarizer, and the load switch, but the setup does include all subsequent RF elements. Over the test range, these forward elements introduce only slight losses, up to a maximum of 0.2dB ripples, which is much less than what the amplifiers produce. It should also be noted that the test signal is monochromatic; we are making the reasonable assumption that the RF system responds to that as it does to a broad-band source.

The two bandpasses are shown in figure 3.3, while the resulting effective bandwidths are tabulated in table 3.1. This bandpass is essentially independent of the attenuator setting (which is the only variable element in

the RF chain), since our Hewlett-Packard step attenuators have differential spectral attenuation deviations of 0.3dB or below, much less than what we are concerned with here.

TABLE 3.1

EFFECTIVE S SYSTEM BANDWIDTH		
Season	1992	1993
$\Delta\nu_{eff}$	5.30 GHz	5.72 GHz




Figure 3.3 Bandpass of the 5.5m telescope's direct-detection receiver. The top panel shows the bandpass of the NRAO system, while the bottom panel shows that of the TRW system (see text).

3.2.4. $1/f$ Noise Behavior

Because of system instabilities, most measurement instruments show increasing noise toward long time scales. In fact, the noise power usually follows a $1/f$ dependence up to the “knee” where white noise takes over. This is one of the fundamental reasons for switching techniques. Does the 5.5m telescope follow this dependence, and does switching alleviate the problem?

The only time this aspect was tested was after the initial direct-detection system was installed with a 3GHz bandwidth, in December 1990. The test simply called for a large number (in this case 200, but 256 would have been better) of short averages. The Dicke switch was set to switch at the fast 2ms period, but the switch was physically disabled. Therefore, data were gathered in switched

mode, but the input data remained as constant as possible, with the load switches pointing to the internal cold loads. The secondary integration time was set to 200ms, and the whole procedure was repeated 10 times.

Figure 3.4 shows the result of this test—first, the total power behavior, then the switched behavior. The total power spectrum follows a $1/f$ law very well; the one drawn has arbitrary normalization. The switched power shows no noise increase toward lower frequencies, which indicates a successful switching technique. In addition, the noise power is




Figure 3.4 $1/f$ noise behavior of the entire system from the internal cold loads backward. The top panel shows the noise spectrum of the total power readings; the bottom panel shows the same data with standard Dicke switching applied. The dotted line indicates a power law slope -1 at arbitrary normalization.

higher than the lowest point reached by the total power spectrum—this is explained by the subtraction operation implicit in the switched power. Even though the data do not sample noise power frequencies above 2.5Hz, we can use the level of the switched power spectrum to conclude that the knee of the total power spectrum probably lies around 2.5–3Hz.

3.2.5. Dicke Switch Isolation

The isolation of a waveguide switch describes the degree to which the power coming into one input port is attenuated when the switch is pointed to the other input port. Typically expressed in dB, we clearly desire a large value. There is no reason that the isolations of the switch directions should be identical, even though it would be desirable.

As described above, there are three switches in the receiver's front end; we can measure the isolations of each. The results for the two load switches are listed in table 3.2 along with the Dicke switch, but

the measurements are explained in the calibration section below. It is almost impossible to measure the load switch isolation when pointed to the sky—fortunately, that parameter is not particularly important, since the internal cold loads are expected to remain at a relatively constant temperature.

The Dicke switch isolation requires a separate measurement, in which the input power to the 'blind' side is varied by a known amount, and the effect in the active side of the switch is observed. This can be accomplished easily by leaving the switch in continuous switching mode and recording the two total power output channels. The advantage of this scheme is also that we measure the effective Dicke switch isolation. If there should be any problems with blanking and other spillover effects during switching, they would enter not only our present test, but also all observations. The results of these measurements are given in table 3.2. It should be pointed out that the errors are calculated from the propagation of the radiometry errors; given my experience over various calibration runs, the systematic errors in these measurements are considerably larger, perhaps about 0.2dB.

TABLE 3.2

WAVEGUIDE SWITCH ISOLATION

Switch, Direction	Symbol	Isolation
Dicke, ANT	$\xi_{D,A}$	$16.547 \pm 0.007\text{dB}$
Dicke, REF	$\xi_{D,R}$	17.710 ± 0.013
ANT Load, Load	$\xi_{L,L}^A$	18.65 ± 0.1
REF Load, Load	$\xi_{L,L}^R$	18.60 ± 0.2

3.2.6. 60Hz Pickup

All electronic systems that are powered by regular AC line power exhibit varying levels of 60Hz interference pickup. In principle, this is not a cause for concern apart from the attendant increase in system noise (which should be avoided if at all possible).

However, we were concerned about possible beating between 60Hz and our sampling frequency. For some integration times (notably multiples of 50ms), this could indeed be a problem. We would therefore like to reduce 60Hz pickup to as low a level as possible. It can be seen from the dashed

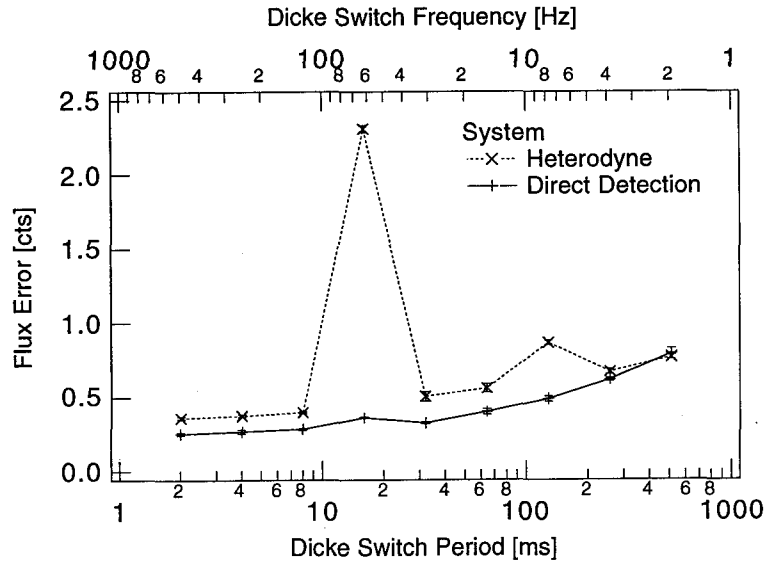


Figure 3.5 The internal scatter in a series of flux measurements as a function of the Dicke switch period. It can be seen that the original heterodyne system is far more susceptible to 60Hz interference than the direct detection system used at present. The $1/f$ rise of the noise level is also apparent; this motivates the use of the fastest available Dicke switch rate.

noticeable 60Hz pickup. We were very pleased to find that the new direct detection system (with which all observations were made) reduces this 60Hz pickup dramatically. This improvement also contains the changes made in the A/D circuitry to minimize AC pickup.

In addition, one can see the rise of the noise level toward lower frequencies; this rise is consistent with the expected $1/f$ behavior. This implies that all observations should be made at the fastest possible switching frequency, which is 500Hz (1ms on each side).

3.2.7. The Beam

A measurement of the beam is fundamental to the experiment. First, it is the most basic test of the system's optical quality. If the system is out of focus, or if there are other aberrations,

tions, these will become **visible** in such a beam map. Second, we can use it to determine the beam's extent on the **sky**, expressed as the beam's full width at half maximum (FWHM). Third, determination of the beam efficiency, a parameter fundamental to the conversion of antenna temperatures to sky temperatures, requires a measurement of the beam's solid angle volume.

In principle, a beam map is easily obtained by measuring a grid of points offset in azimuth and zenith angle from a bright point source. Yet good beam maps are difficult measurements, especially probing the beam's outer areas. Gain and switched power drifts can easily introduce considerable artifacts into the beam's apparent shape. For this reason, I measured the beam map in a 23×23 grid with a spacing of $1'.2$, scanning in azimuthal lines, but returning to the peak every 4–5 grid points to adjust the pointing offsets and to measure the system gain and the switched power offset. In addition, a proper map that probes the beam's wings sufficiently takes about 6 hours of observing time in good weather.




Figure 3.6 Contour map of the effective Dicke switched beam of the 5.5m telescope. This map shows the response function on the sky. The ANT beam is on the left, while the REF beam, with negative contours, is on the right. Contour levels are $0, \pm 10, \dots, \pm 90\%$ of the peak. The ragged vertical contour between the two beams (at the zero level) is caused by the overlap of the two originally independent beam maps.

Using Jupiter, such a beam map was measured 4 independent times on the ANT side of the telescope; there was time for only one corresponding REF beam map of sufficient quality.

The three usable ANT maps and the REF map were combined to form the composite (switched) beam map shown in figure 3.6.

It can be seen that the beams are well-behaved, symmetrical, and very similar to each other. Fitting a two-dimensional Gaussian to each map, we find that the two beams follow that shape very well; rms deviations from a Gaussian fit are less than 2% of the maximum, and less than 5% at the peak. The resulting FWHM is 7'.03, and the ANT beam's volume (which is better determined than the REF beam's) is 5.16×10^{-6} ster. The error in that value is estimated at 3%, or 0.15×10^{-6} ster.

3.2.8. Pointing

A telescope's nominal pointing direction does not correspond to its true pointing direction in the sky to sufficient accuracy. It is not only necessary to establish the zero points of the encoders, but to compensate for effects like the non-intersection of axes and their non-orthogonality, for tilts in the system, etc. Usually, this is achieved by looking at a large number of known radio sources, well spread over the sky, to peak up on these sources with the telescope's radiometer, and to record the necessary pointing offsets to bring the center of the telescope's beam directly over the sources, as a function of azimuth and zenith angle. From this, then, one can derive a model for the telescope's pointing corrections. This model becomes part of the telescope control program and is applied automatically to all requested positions.

Over the years, the following pointing model has been developed. The total azimuth and zenith angle corrections, C_A and C_Z , respectively, are calculated from the nominal azimuth A and zenith angle Z and the pointing constants a_{1K8} and z_{1K4} using the following relations (all in arc minutes):

$$C_A = a_1 + a_2/\sin Z \\ + (a_3 \sin A + a_4 \cos A + a_5) \cot Z \\ + a_6 \sin A + a_7 \cos A + a_8 \sin 4Z/\sin Z ,$$

$$C_Z = z_1 + z_2 \sin Z + z_3 \cos A + z_4 \sin A .$$

The main problem with these pointing equations is that their terms are not orthogonal. In some cases, in fact, two seemingly independent terms may be the coefficient for one physi-

cal term, such as a wobble term. This and the rather unconstrained fitting software need major refurbishment, but this was not of critical importance to the thesis' scientific goals.

After each physical change in the drive system (*e.g.*, encoder removal and re-installation), the pointing constants had to be determined again. While a significant asset to our program observations, the small telescope's low point source gain poses significant problems for pointing tests. Only two standard radio sources are bright enough to serve as pointing sources: DR21 and 3C84. Among the planets, only Venus and Jupiter, and sometimes Mars, are bright enough to be used for pointing purposes. In general, the pointing data from Venus and Jupiter have much better quality than those from the radio sources—for this reason it usually takes upward of three full days in good weather to gather enough data for a satisfactory pointing solution.

Typically, the final rms fit in a pointing solution is about 23"—about one-twentieth of the beam's FWHM. During the last pointing run before the 1993 season, however, the fit's rms remained above 30". It was then found that the pointing residuals for the planets varied from day to day by as much as a half arcminute, while they stayed constant for the radio objects. Meanwhile, this problem has been traced to a synchronization problem in the planetary pointing code, and has been corrected for future observations. Fortunately, all deviations are restricted to zenith angles greater than 40°, a region in which only a small fraction of the observations were taken. In addition, the pointing error is small (about 30"), compared to the beam size of 7 arc minutes.

Using pointing data that were taken with the ANT and the REF beams independently, the effective separation of the beams on the sky can be measured. We found a consistent separation of 22'.16 for the two beams—this parameter has been adopted as our nominal beam separation. This is consistent with the design separation of 3 beamwidths.

3.2.9. Zero Levels

Are there any low-level biases from the backend that could have a systematic effect on our results? If, for example, the system's zero levels showed any systematic bias, then the quality of the entire system's results would be compromised. Fortunately, zero level tests are easily performed, and there are many of them. Two features are highlighted here.

First, the zero level has been observed to drift over a total range of about two thousand counts. Sometimes, the drift can be rapid—spanning as much as 300 counts over 12 hours. They originate from before the backend, since they disappear when the RF inputs are disconnected; a likely source is the double detector assembly, whose differential output can easily drift by several mV if the detectors' physical temperature changes by 1K.

The biasing effect of such drifts on the switched output is minimal, however. A 300-Count drift over 12 hours translates into a $1.4\mu\text{Count}$ offset with 2ms Dicke switching, our standard observing configuration. At standard observing power levels, this translates into a 20nK offset, surely not detectable in even the longest integrations.

The power spectrum of a long zero level test with 5-second sampling reveals two spikes corresponding to periods of 25.73 and 51.36 minutes, apart from the inevitable $1/f$ rise. No instrumental explanation has yet been found to these frequencies, but they are slow enough to be eliminated by Dicke switching.

A second, potentially more troubling, feature of the zero level is that its switched value is consistently negative by about 0.1 counts. This switched value offset is persistent and constant; it can be seen at the same level, with the same sign, in data separated by a year. Even the replacement of most of the backend's analog circuitry did not remove this offset. This offset is baffling: it comes from a section in the receiver that is by design not aware of the direction in which the Dicke switch is pointing.

Fortunately, this offset does not affect our data at a measurable level. In a 5-hour test, 1664 zero level integrations were made. The switched version of these measurements shows the offset at a level of $-123\pm 2\mu\text{K}$ at equivalent observing power levels. An offset in a switched quantity does not automatically translate into an offset in the flux measurement. Indeed, the offset disappears when the Flux algorithm is applied to these switched data. With a simulated flux segment duration of 10s, an offset of $-9.1\pm 6.1\mu\text{K}$ remains, while with a duration of 20s per flux segment the offset is $-5.8\pm 5.7\mu\text{K}$. Neither value is statistically significant, and in any case small enough to be disregarded. Any remaining offset here would be removed by the third differencing step during observations.

3.2.10. System Stability and the Length of Flux Measurements

To decide on an optimal observing technique, we need to know how the length of a flux measurement affects the stability of the data. This stability is affected by all levels of the observations—not only the instrument but also the atmosphere. The trade-off is clear: on one hand, shorter switching times are better, since they reduce the time base of fluctuations. On the other hand, shorter position switching times within a flux measurement dramatically reduce the observing efficiency, since more telescope moves are required.

Using the NRAO system, I performed a test to see if there was any significant difference between the internal scatter of a longer-term average of Flux measurements, depending on whether individual segments had a length of 10s or 20s. The noise in individual Flux measurements is within 10% of the thermal noise expectation, independent of the segment integration time. The question is rather whether the scatter in a long average of such flux measurements corresponds to the thermal expectation or whether the long-term noise does not decrease with the square root of time. This is equivalent to measuring the dynamic system temperature on longer time scales with a clear sky (see section 3.2.2).

I found that in both cases the dynamic system temperature was substantially larger than the nominal system temperature on time scales longer than 10–20 minutes, but no significant difference between 10s and 20s fluxes was seen. In the case of 10s fluxes, the dynamic system temperature from the average of 46 measurements spanning 53 minutes was 2.2 times larger than the nominal value. In the case of 20s fluxes, the dynamic system temperature from the average of 22 measurements spanning 40 minutes was 2.1 times larger. Allowing for the different total averaging times, these two numbers are certainly equivalent, but disappointingly large.

This increase in the dynamic system temperature from nearly its nominal value on time scales of minutes to more than twice that amount on time scales of an hour is disappointing, but very difficult to characterize, even though we suspect atmospheric fluctuations as its main cause.

It would certainly be useful to perform a more detailed study of the entire system's noise performance. However, these tests are very difficult to perform. Since they are meant to describe the noise performance of the entire system, they must be performed while looking at the sky, and require the kind of excellent observing weather that is not easily given to

such tests. However, it appears unlikely that the dynamic system temperature increases much at even longer time scales: the total SZ result from the 1992 season has an internal scatter 2.5 times the thermally expected noise, only slightly larger than the ratio observed in these tests.

In principle, it would be possible to use the internal cold loads to determine if the receiver system introduced this additional long-term noise, but it has become apparent that long-term stability cannot be studied well with the internal cold loads. The problem is simply that the two loads' emissions are only approximately physically correlated, in the sense that they reside in the same dewar and that they are cooled by the same refrigerator. However, minute temperature changes, which the system measures very well, are not correlated between the two sides, so that we are unable to gain information on the system's stability because the test sources are not sufficiently stable.

Given the receiver's performance in other aspects, I find it unlikely that the rise in the dynamic system temperature comes from some long-term system instability. Even if such instabilities existed behind the Dicke switch, they would be very well subtracted by the constant switching operation. In my opinion, atmospheric instabilities account for the bulk of this problem.

3.2.11. Ground Pickup

Shortly after the start of our first HEMT-based observations in December 1990, some puzzling results were obtained. In observing 12 intrinsic microwave background fields around the north celestial pole in excellent weather, the data values drifted by several mK away from zero. No real sky source could possibly account for such a large signal, and receiver problems appeared an unlikely cause. After some time, we found that the dubious temporal profile of each of the 12 fields was repeated when it was re-observed after 24 hours, and we started suspecting ground pickup as a cause of the problem.

Initial attempts to cover the hole in the dish through which the TV camera aligns with the telescope's radio axis did nothing to improve things, even though such a major radiation leak should remain covered (as mentioned in section 3.1.2, it remained covered for subsequent observations). Further attempts sought the problem in a possible mis-illumination

through the feeds, but an absorbent annulus mounted on the secondary to reduce its reflective diameter failed to improve the results.

Finally, crinkled aluminum foil was wrapped around the telescope's feed legs, and the signature of the ground pickup problem changed dramatically. It became obvious that the massive feed leg arrangement, which gave enormous mechanical stability to the position of the secondary reflector, was angled such that ground radiation could enter the parallel optical beam in a single reflection at virtually all zenith angles! This support leg structure, a quadrupod with split and flared legs, used 1.625-inch tubing for an effective 8 legs that were tilted by about 45° to the telescope's line of sight. Strong total ground pickup would result from this arrangement, and two levels of differencing in the form of double-switched flux measurements (see section 4.2) were unable to reduce this bias to an acceptable level.

The total ground pickup, i.e., the total increase in the system temperature due to reflected ground radiation, was about 27K, too much to tolerate even if no systematic biases would result. After the two levels of differencing, the resulting differential ground pickup would drop to several mK, but as a strong function of zenith angle. To appraise the effect properly, we conducted a great many series of full flux measurements with the telescope pointed to a fixed altaz position—not tracking a source over the sky. Instead, the sky then drifted behind the pointing position, providing a reasonable average sky temperature. These flux measurements were then repeated at various zenith angles, so that the variation of the effect with position could be appreciated and so that a good estimate of the effect's peak value could be made.

The peak differential ground pickup in these measurements was almost 5mK, clearly too much for a system that aims for residual systematic biases of $10\mu\text{K}$. It became clear that the solution for this problem was to replace the entire secondary support leg structure, but this meant that the 1990–91 observing season had to be abandoned entirely.

The ideal solution to this problem would be to replace the entire telescope with a clear-aperture design that would consequently sport extremely low sidelobe levels. However much desirable, this solution was beyond our means, and we decided to design a secondary support leg structure with minimum total and differential ground pickup. The details of this design are explained in a paper accepted for publication in the Proceedings of the IEEE. A preprint of that paper appears in appendix B.

The new design was very successful. The total ground pickup with the new design dropped to 9K, a respectable value for a high-frequency radio telescope. But more importantly, the new design gave a peak differential ground pickup of $150\mu\text{K}$, a factor of 30 below the previous level! At most zenith angles, this value was even lower, in the range of several tens of μK . Of course, testing of the new structure had to be stopped at some point, since bias tests to a level of tens of μK are extremely time-consuming and require excellent observing weather. A good illustration of the change is given by figure 3.7, which also appears in our IEEE paper. Because it summarizes the improvement very well, it is reproduced here.




Figure 3.7 Differential ground pickup (see text) from the 5.5m telescope. Part (a) shows the differential ground pickup at a series of zenith angles using the original secondary support leg structure; part (b) shows similar data taken after replacement of the support structure with a low-pickup design.

Still, a level of tens of μK of ground pickup bias is clearly too much for a system that aims for residual systematic biases of $10\mu\text{K}$. However, it appears plausible that a further level of differencing, if correctly arranged, could succeed in subtracting this bias from our data. Such a scheme is introduced in the section on observing techniques below.

Parts of the new design were based on sketches provided by J. Welch, and significant help came from the antenna group at JPL, in particular D. Bathker, P. Cramer, and R. Levy, who suggested the addition of a vertex plate to the secondary and who performed the mechanical rigidity analysis. The design called for minimal optical obstruction, while a high level of positional accuracy of the secondary is maintained.

The vertex plate is a small insert into the central 3-inch portion of the secondary reflector, flaring out slightly from the surface to form a point at the center. It helps redirect the bulk of the illumination that would normally hit the center of the primary, where it would either

scatter from the protruding front of the receiver or be obstructed by the secondary itself. In the original design, a whole 6% of the illumination occurred in this obstructed region because of the centrally weighted illumination pattern. Since the eventual disposition of those rays is not clear, such a large fraction is clearly undesirable. With the vertex plate in place, however, this portion could be reduced below 1.5%. The effective impact of the vertex plate on the differential ground pickup cannot be easily judged, however, since it was installed at the same time as the new support leg design.

The new system was in place for observations in the fall of 1991.

3.2.12. Non-linearity

It was discovered only in early 1993 that the radiometry system exhibits some amount of non-linearity. It is a small amount (about 5–9% over a 5dB range), but it can affect the data calibration by an equivalent amount, and it can have a large (up to 30%) effect on the system temperature that we derive from our calibration measurements. It is important to characterize this non-linearity not because of the way we observe (we stay at about the same power level during all aspects of the observations), but rather because the calibration procedure assumes a linear system over a dynamic range of 5dB. This assumption during calibration could be eliminated if we had a precision-calibrated attenuator with sufficient resolution. In that case, we could simply take the hot load reading at one attenuator setting and then decrease the attenuation by a known amount to bring the power level into a similar range to measure the cold load. Unfortunately, such an attenuator would be prohibitively expensive, given the level of precision required.

Instead, we can use two avenues to deal with this problem. One is to eliminate the non-linearity, and the other is to map it and apply it as a correction to all total power data. Since we are not quite sure where the non-linearity occurs (in fact, it may even come from two independent sources), the first option was not achievable in time for the start of observations. A careful mapping of the non-linearity, however, is a good alternative as well, and has been performed with the 1993 season's system.

Unfortunately, there was a much more unpleasant implication: if there is non-linearity in the TRW system, and since its source is unlikely to be in the first-stage amplifier at low power levels, we probably had a non-linear system during the previous season without

realizing it. Fortunately, the NRAO amplifier was still available, and tests were performed in June 1993 with the old amplifier, a configuration very similar to the one used in that season, except for the changed integrator and the A/D components that were replaced in the summer of 1992. Because of the stringent linearity specifications of these components, however, I regard them as a very unlikely source of the observed non-linearity. As described below, the main portion of the non-linearity must come from before the RF attenuators; this effectively excludes the backend electronics from suspicion. More likely, the observed non-linearity is a combination of saturation non-linearity from the last Avantek RF amplifier and from the differential detector system.

Formalism

To describe the non-linearity, a simple formalism was developed. Non-linearity is another way of saying that the observed output power and the true input power are related by some non-constant mapping $\eta(P_o)$, so that

$$P_o \equiv P_T \eta(P_o).$$

Non-linearity is best measured differentially in terms of (small) power changes, such as the CAL, where the observed CAL is related to the 'true' CAL by

$$C_o(P_o) \equiv C_T \lambda(P_o).$$

In the absence of an analytic expectation for $\lambda(P_o)$, we use the simple fit of

$$\lambda(P_o) = a + bP_o.$$

The differential non-linearity is related to the total non-linearity by

$$\frac{\partial \eta}{\partial P_o} = \frac{\eta}{P_o} - \frac{\eta^2}{\lambda P_o},$$

which for our choice of $\lambda(P_o)$ and the boundary condition of $\eta(0) = 1$ has the solution

$$\eta(P_o) = \frac{\lambda(P_o) - 1}{\ln \lambda(P_o)}.$$

Once the function $\lambda(P_o)$ is measured, all total power measurements as well as all differential measurements can be corrected. Since a is a free parameter in the λ -law, the parameters are normalized so that $a = 1$.

Skydip Test

Initially, the non-linearity was discovered by monitoring the physical temperature and the emission power of the internal cold loads while heating the loads with their built-in heaters. During this procedure, the value of the CAL was monitored to correct for variable system gain. As the internal cold loads heated up and the input power increased accordingly, the CAL value did not remain constant, but rather dipped by about 3.5% in unison with the input power increase.

Even though this test served to discover the problem initially, it could not be used to map the effect, since it only covered a dynamic range of 2 in input power. Another test was needed in which the input power could be varied over a larger dynamic range without changing the RF attenuation. This was achieved by increasing the RF input power simply by increasing the amount of emitting atmosphere in the telescope's beam, while using the CAL at each input power to measure the non-linearity. This can be done at a variety of attenuations, and independent λ -laws can be fitted to the data for each. We can then compare the λ -laws at different attenuations to see if there are significant and systematic trends in the fit parameters that correspond to the changing attenuation.




Figure 3.8. Measurements of the non-linearity parameter b as a function of attenuator setting with both systems. The power levels of the TRW system lie about 4dB higher than those of the NRAO system, and the non-linearities in the two systems are consistent with each other (see text). The dotted lines denote fits to the data and are described in the text. The non-linearity is described by $\lambda = a + bP_0$, normalized so that $a=1$.

The normalized fit parameters b are plotted in figure 3.8 as a function of attenuation for each system. In both systems, a correlation of b vs. attenuation is visible—this trend is entirely expected, since higher attenuations compress the observed power scale. If all the

non-linearity comes from before the attenuator, the parameters b derived at two attenuations should relate to each other as

$$b_1 = b_2 A_{1,2} = 10^{(dB_1 - dB_2)/10}.$$

By fitting to the data in figure 3.8, we can determine what fraction γ of the non-linearity comes from before the attenuator. In addition, the fit will also yield the parameter b to be used at each attenuation. The fits worked remarkably well, especially for the TRW system, as can be seen by the dotted lines in figure 3.8.

We can also use these fits to a partial attenuator dependence to compare the two systems. Since only first-stage amplifiers have changed between them, and since those are extremely unlikely to contribute to the non-linearity, we would expect to see some

similarity between the two systems. This is illustrated in table 3.3. At the two overlapping attenuator settings, the two systems are compatible with each other.

For the purpose of non-linearity correction, we will simply make use of the fact that all observations take place at the same attenuator setting (as do the calibrations), and use the non-linearity parameters resulting from the fits at various attenuations. Table 3.4 summarizes the non-linearity parameters to be used in various cases.

TABLE 3.3

PRE-ATTENUATOR NON-LINEARITY FRACTION			
System	Attenuation	γ	$b_{tot} [\times 10^{-6}]$
NRAO, 1992	2dB	9.6±2.0%	-1.423±0.048
	6	21.1±3.9	-1.630±0.030
	8	29.8±4.9	-1.831±0.048
TRW, 1993	6dB	23.2±1.5%	-1.678±0.029
	8	32.4±1.8	-1.905±0.023
	12	54.6±2.1	-2.838±0.047

TABLE 3.4

FULL SYSTEM NON-LINEARITY PARAMETERS			
Domain	Attenuation	$b [\times 10^{-6}]$	
NRAO system, 1992	observations	2dB	-1.423 ± 0.048
	system calibration	8	-1.831 ± 0.048
TRW system, 1993	observations	6	-1.678 ± 0.029
	system calibration	12	-2.838 ± 0.047

3.2.13. CAL Diodes

Our entire calibration relies on the performance and stability of the calibration noise diodes. This necessitates several tests to characterize them properly. The first question is about

their short-term stability. We can measure this easily by comparing the value of the CAL's power as a function of delay after the diode is turned on (parameter I in section 3.2.1), as well as a function of the length of the Cal procedure's segment integration time S .

Short-Term Stability

In the first test, where the value is measured as a function of the idle parameter I , no dependence could be found, except for the shortest idle times of 1s (the CALs indeed take a few hundred milliseconds to turn on). For parameters larger than that, the CAL value in either channel remains unaffected to less than 0.2%. For the next test, involving the length of the Cal integration, a clearer signature is discernible (see figure 3.9)

Here, there is a clear increase of the measured value with longer integration times. This is somewhat puzzling, since the CALs do turn on quickly initially. Perhaps the CALs warm up while on, which may increase their power output. By the same token, one would expect to see this effect in the variation with I , but it is not seen there.

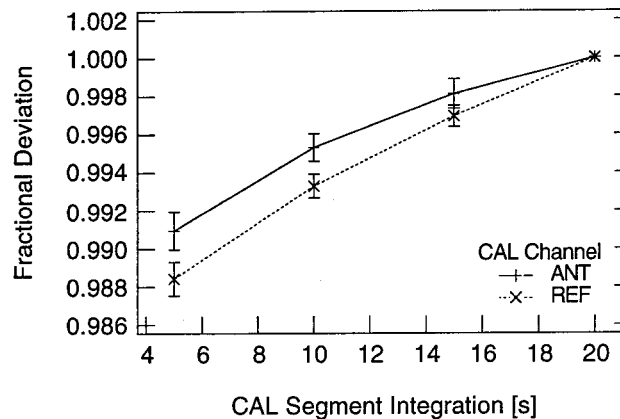


Figure 3.9 Change of the value of the calibration diode in Cal procedures of varying segment integration times.

Since the system calibration uses segment integration times of 10s, while the CAL during observations is used with an integration time of 20s, a correction factor of 0.5% will be used below to increase the temperature equivalent of the CAL in the ANT channel as measured by the system calibration.

Switched vs. Total Power

There is also a distinction between measuring the CAL as a change in the total power versus as a change in the switched power, since imperfect isolation between the two channels will decrease the measured switched signal. Tests to measure the Dicke switch isolation with

this method are not nearly as accurate as the direct tests described in section 3.2.5, but the results must be consistent. This is indeed the case. By comparing the switched and the total-power calibration values (while the backend remains in switching mode, so that the blanking remains constant), the Dicke switch isolation for the ANT and the REF channel are -15.9 ± 0.7 dB and -17.7 ± 1.2 dB, respectively. This is entirely consistent with the more definitive values above.

The distinction of switched and total power CALs is important to keep in mind in deriving and applying the system calibration below. The two CALs are related by

$$T_{Cal,S} = T_{Cal,TP}(1 - \xi_D),$$

where ξ_D is the Dicke switch isolation of the opposite channel. Similarly, a real sky signal in one channel is diluted by the same imperfect isolation

$$T_{Sky,S} = T_{Sky,True}(1 - \xi_D).$$

While in switching mode, we measure the sky and the CAL in essentially the same fashion (P_S is the switched power observed)

$$T_{Sky,S} = \frac{P_{Sky,S} T_{Cal,S}}{P_{Cal,S}}.$$

From this follows that the true sky temperature can be calculated as

$$T_{Sky,True} = \frac{P_{Sky,S}}{P_{Cal,S}} \frac{T_{Cal,S}}{(1 - \xi_D)} = \frac{P_{Sky,S}}{P_{Cal,S}} T_{Cal,TP}.$$

In other words, if we observe both the sky and the CAL in switched mode, and use the absolute temperature of the total power CAL, no further correction for the Dicke switch isolation is required.

Long-Term Stability

We are making one other assumption when using the CAL—namely, that its power output remains constant over a long period of time. This question is very difficult to settle, and we do not yet have a definitive answer for it. The difficulty lies in the fact that the CAL must be measured externally, in independent, but regular system calibrations, which are frequent and consistent enough to be tied to each other. As explained below, such a facility for system calibration measurements is a difficult proposition, and has only recently been

solved for this instrument. As a result, only sporadic and somewhat inconsistent measurements of the CAL have been available over the years.

From this, we can see that the long-term CAL stability is not as good as expected. For example, the measured CAL value for the ANT channel was 470mK in March 1993, while it had apparently climbed to 485mK by the middle of April, all the time with the exact same system setup. This 3% change is not explained, nor do we even understand how significant it is. On the REF side, the CAL has changed more: from 519mK to 480mK in the same time period. While the significance of the magnitude of the change is unclear, there is no doubt that one or both of the CALs are not constant in the long term: we can measure the ratio of the two CALs during the same system calibration to high significance. This ratio changes from 0.983 ± 0.003 to 1.040 ± 0.004 in the same time period, which is a change of 0.057 ± 0.005 , a significant effect. No doubt more effort has to be invested tracking down the source of this change.

3.2.14. Outstanding Problems

A hardware project rarely has everything working perfectly nor does it ever achieve a state from which it cannot be improved further. The 5.5m telescope is no exception to this, but I believe that it is entirely adequate for its task. Especially considering the enormous improvement in system sensitivity, I also believe that we have succeeded in putting the main effort into those components of the system that truly count for our observations.

Still, many things need improvement and change. For one, unacceptable amounts of observing time have been lost to drive and encoder failures; this has improved since the drive motors were rebuilt in the summer of 1992. Yet encoder problems and problems with the vintage drive controller persist. Because they usually occur around nightfall, when the outside temperature decreases rapidly, the engineering staff has difficulty tracking them down.

Another vital improvement is a better pointing model and better pointing software for the telescope. A new algorithm and associated software now exists to analyze pointing data, and we are planning to use substantial of observing time during the end of the summer to concentrate on pointing exclusively. This should solve the pointing inconsistencies we have seen so far.

Another point of concern is the long-term stability of the CAL diode, especially as expressed by the changing ratio of ANT/REF. We will monitor this ratio during observation in the future to see whether such changes are of a diurnal or of a secular nature. In addition, completion of the cold load calibration box (described in section 3.3.2) will allow much more frequent system calibrations, which should be performed at least once a month during the observing season.

Even though it is not vital to the quality of the observations, it would be desirable to remove the non-linearity from the system. Since it most likely comes in part from the third RF amplifier and in part from the differential detector system, its cause should be identifiable and perhaps easily remedied.

3.3. CALIBRATION

3.3.1. System Temperature Measurements

System temperature measurements aim to determine the noise temperature contributions by all the experiment's components, including the radio source under question. This total system temperature T_{sys} includes the noise contribution due to the receiver, but also ground spillover, atmospheric emission, and the microwave background. Once a calibration has been established, the total system temperature equals the total power seen by the radiometer, and is therefore easily measured. On the other hand, it is not so easy to apportion the noise power to various components. The contribution to the system temperature from the receiver itself can most easily be measured during the system's calibration.

The basic calibration principle for a total power radiometer is to measure two black bodies of known emission temperature as well as the instrument's zero level, and to derive from this the scale (*i.e.*, the temperature calibration) and the offset (the receiver temperature) of the correspondence. With a complex receiver structure, the receiver temperature T_{rx} consists of several contributions as seen in figure 3.10. Significant emitters within the receiver are the horns (T_{Horn}), the internal cold loads (T_{IL}), the CAL noise injection (T_{Cal}), and the first-stage HEMT amplifier. The other elements, such as the polarizers, Dicke switches, cross-guide couplers, and waveguide flanges are assumed to first order to contribute no noise. Indeed, some of these lesser elements are simply part of others, for example, the

polarizer emission is considered part of the horn emission, and is not separable with this technique.

Other parameters of interest are the Dicke and load switch isolations: the attenuation of radiation seen from one input port of the switch when it is pointed to the other port. The load switch isolation (when switched to the internal load) can be derived from the system calibration procedure, but the Dicke switch isolation must be known independently (see section 3.2.5). The load switch isolation is denoted ξ_L with the subscript L or S if the switch is pointing to the internal cold load or to the sky, respectively. The Dicke switch isolation is denoted ξ_D , with the subscript A or R depending on the position of the switch.

Finally, there is the gain parameter g which carries the conversion from the instrumental counts to temperatures; this parameter is valid only within a calibration run, since we would expect the instrument's gain to change (this is why we use the CAL injection).

For each external load (hot and cold), three measurements are taken as well as a zero level determination. Below, we assume that this zero level has been subtracted from the measurements. In the ideal case, the three measurements are taken on both sides simultaneously. First, we make an integration on the external load, then an integration on the internal cold load (with the same external load in place), and finally an integration on the internal cold load with the CAL injection on. Experimentally, the first integration is then repeated to provide a stability estimate. These three integrations are called A, B, and C; they are usually performed with the hot, then twice with the cold, and finally with the hot load again. Subscripts C and H denote integrations with the cold and the hot external loads, $T_{L,C}$ and $T_{L,H}$, respectively.

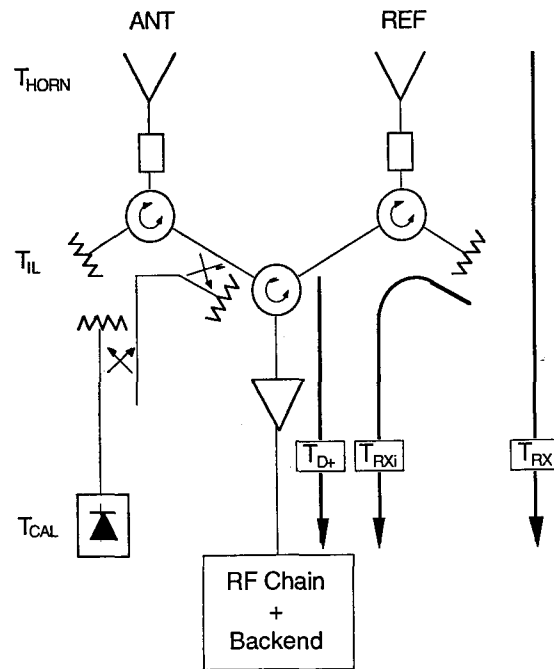


Figure 3.10 Schematic of the receiver's front end, showing the paths measured during a system calibration.

The three integrations can be expressed as follows:

$$\begin{aligned} A_k^A &= g^A (T_{Hom}^A + T_{D+}^A + T_{L,k} + \xi_{L,S}^A T_{IL}^A + \xi_{D,A} T_{L,k}) \\ B_k^A &= g^A [T_{D+}^A + T_{IL}^A + \xi_{L,L}^A (T_{Hom}^A + T_{L,k}) + \xi_{D,A} T_{IL}^R] \\ C_k^A &= B_k^A + g^A T_{Cal}^A, \end{aligned}$$

where $k = \{H, C\}$, and where ANT and REF are exchanged by switching the indices A and R. From these measurements, we can then calculate the y-factor

$$y^A \equiv \frac{A_H^A}{A_C^A},$$

which is used to derive the receiver temperature $T_{Rx} \equiv T_{Hom} + T_{D+} + \xi_{LS} T_{IL}$ and the gain as

$$\begin{aligned} T_{Rx}^A &= (1 + \xi_{D,A}) \frac{T_{L,H} - y^A T_{L,C}}{y^A - 1} \\ g^A &= \frac{A_H^A}{T_{Rx}^A + (1 + \xi_{D,A}) T_{L,H}}. \end{aligned}$$

Further, the internal receiver temperature $T_{Rx,I} \equiv T_{D+} + T_{IL} + \xi_{L,L} T_{Hom}$ can be calculated as

$$\begin{aligned} y_B^A &\equiv \frac{B_H^A}{B_C^A} \\ T_{Rx,I}^A &= \frac{(\xi_{L,L}^A T_{L,H} + \xi_{D,A} T_{IL}^R) - y_B^A (\xi_{L,L}^A T_{L,C} + \xi_{D,A} T_{IL}^R)}{y_B^A - 1}, \end{aligned}$$

while we can calculate the load switch isolation in the load position as

$$\xi_{L,L}^A = \frac{B_H^A - B_C^A}{g^A (T_{L,H} + T_{L,C})}.$$

The CAL noise temperatures follow easily from these two independent calculations:

$$\begin{aligned} T_{Cal,H}^A &= \frac{1}{g^A} (C_H^A - B_H^A) \\ T_{Cal,C}^A &= \frac{1}{g^A} (C_C^A - B_C^A), \end{aligned}$$

even though this is not a very accurate method for measuring them. It is much more accurate to use full Cal procedures for this purpose, which becomes possible if the external loads are stable (see below). Combining these with the derived value for the gain g , the noise temperature of the CAL is obtained.

Unfortunately, it is not possible to distinguish between horn temperature and internal cold load temperature with this procedure. However, physical thermometry is available for the internal cold loads. Assuming that the calibration is correct and that physical temperature indeed corresponds to emission temperature, we can calculate T_{Horn} as

$$T_{Horn}^A = T_{IL}^A - \frac{B_C^A / g^A - T_{Rx}^A - \xi_{L,L}^A T_{L,C} - \xi_{D,A} T_{IL}^R}{1 - \xi_{L,L}^A}.$$

Error propagation of these quantities shows a very strong dependence on the external cold load temperature—its value should therefore be known to within 0.5K.

3.3.2. External Cold Load

To carry out the above system calibration procedure, we must possess reliable external black-body emitters. Also called loads (because of their desired property of not reflecting any radiation), they usually involve a radio-frequency absorbent material. In our case, we relied on various Eccosorb™ foam absorbers manufactured by Emerson and Cuming.

Among loads, a hot load is the simplest to obtain: take a piece of Eccosorb and hold it in front of the feeds. Depending on the thickness of the absorber, a reflective back plate must be used to prevent transmissions through the foam. Our standard hot load consists of a round patch of Eccosorb AN-73 absorber that is attached to an aluminum plate. According to the specifications, reflections from this absorber should be suppressed at the 50dB level. All that remains is to monitor the load's physical temperature with a thermometer. In many cases it was sufficient to use the thermometry information from the site weather station to provide the hot load temperature.

The cold load is another matter entirely. The demands on a liquid nitrogen cold load are much more stringent than those for a hot one: the derived system temperature is about 6.5 times more sensitive to temperature errors in the cold load than to errors in the hot load.

Cold loads usually consist of an Eccosorb piece cooled by liquid nitrogen (LN₂). Commonly, cold load cones are employed—copper cones lined with Eccosorb that are first submerged in a bath of LN₂ and then lifted out and held in front of the feed of interest. This is a problematic procedure for several reasons. First, it is fairly dangerous. Liquid nitrogen is dripping from the load while measurements are made. Second, it is not stable:

the LN₂ boils off within 20–30 seconds, and the emission temperature begins to rise as a result. Third, the cone is usually too small to cover both feeds. Consequently, the second feed looks at some unspecified temperature, which is seen in the calibration channel because of the imperfect isolation of the Dicke switch. This is especially serious in receivers that are not equipped with internal cold loads. Further, the cold load cone warms up starting from the rim. Since the cone used for the 5.5m telescope just barely fits over the feed, the effective emission temperature is always in question. A larger cone load would require a bigger and bulkier dewar, with all the attendant problems of weight and safety. But most problematically, because of the porous nature of the absorber and because of its exposure to air during the procedure, the effective emission temperature is never clearly defined.

A much easier design is an absorber that remains submerged in a bath of liquid nitrogen. However, several requirements must be met for this to work. First, the LN₂ container must be well insulated. Not only does this keep LN₂ from boiling off, but it prevents condensation from forming on the container, which would thoroughly contaminate all test measurements. Secondly, the container's walls must be transparent to radio waves. Third, reflections from the wall/nitrogen interface must be prevented.

This last problem is particularly important, since the index of refraction of liquid nitrogen is about 1.20 at radio frequencies (NBS Circular 514), leading to reflections of about 1% of the incident power. Since the reflections would include radiation emanating from the internal cold loads through the load switch, some of the reflected radiation would be substantially colder than the calibration temperature. Without doubt, some radiation from the outside would also be reflected into the feed by this interface, leading to an unspecified increase in the calibration temperature. Clearly, without some effort to prevent reflections, the effective radiation temperature of the cold load would not be known within 1K.

To reduce reflections from the air/nitrogen interface, an anti-reflection coating is desirable. Since the refractive index of most foams is very close to that of air, and since foams with higher refractive indices typically have unacceptably large absorption at 1cm wavelengths, a simple foam layering technique was not possible. Instead, we decided to drill a hole pattern into the inside foam surface to achieve a median refractive index by allowing the foam's material to 'mix' with the nitrogen at the interface. Because of the wavelength of 9mm, small holes were required, which in turn affected our search for a suitable foam material. Not only does the foam have to provide low radio emission and very good ther-

mal insulation, coupled with minimal thermal shrinkage to reduce the wear on the container by repeated thermal cycling, but the material also had to be strong or flexible enough, with small enough cells, to make such a hole pattern feasible.

A large number of foams was investigated for this purpose. Low-density polyethylene foam, a somewhat translucent large-cell material, was the winner from the radio emissions standpoint, having emission that was measured at well below 0.1K for a sheet with a thickness of 2 inches. However, it leaks liquid and performs badly in all other requirements. Styrofoam, an extremely rugged and easily machineable foam containing a blue fire-retarding compound, fulfills almost all requirements except that of low radio emission. Finally, we found a cross-linked polyethylene foam of much smaller cell size than the above-mentioned polyethylene foams with good mechanical properties and moderate emission of a few K for a 2-inch sheet at a physical temperature of 300K.

The anti-reflection hole pattern is designed to approximate the refractive index in the boundary layer that is required to cancel reflections. The desired median refractive index can be calculated by requiring exact E-vector cancellation at the interfaces, and is achieved by drilling out an appropriate volume from the foam to a depth of $\lambda/4$ at the intermediate refractive index. Since the hole pattern is on the inside of the foam container, the holes are filled with liquid nitrogen. To achieve the correct intermediate refractive index of 1.105, the ratio of drilled-out to filled volume should be 0.818. We decided that a number of 4 holes per wavelength in the transverse direction was acceptable, allowing us to use an end mill with a diameter of 0.0625 inches. The depth of the holes had to be 0.0835 inches to take account of the intermediate refractive index. The resulting hexagonal hole pattern of about 7500 holes covers the 5×10-inch area of the illumination of the two horns.

To reduce the radio emission from this section of foam (at an unknown effective emission temperature), we machined a pocket on the outside of the illumination patch, so that a thickness of 3/8 inches of foam was left, including the anti-reflection pattern, and to provide a seal for the liquid nitrogen container. This pocket was filled by a very low-density open-cell polyethylene foam plug to provide mechanical stability against bulging out due to the pressure from the liquid contents, and to provide a thermal insulator to prevent condensation from forming on the outside of the thinned foam section.

The absorber immersed in the liquid nitrogen is Emerson and Cuming CV-3 with a pyramid pattern and increased thickness so that a reflective back plate is no longer required.

Using the above construction, we estimate that the emission from the thinned foam sheet is about 0.5K, which must be added to the emission temperature of the cold load. Given a boiling point of liquid nitrogen at barometric pressures of 880mBar at 4000 feet altitude of 76.1K (see HCP), an effective cold load emission temperature of 76.6K should be used. The calibration section below discusses the results from this cold load container.

3.3.3. Calibration

System calibrations were performed at several times during the period covered by this experiment. Unfortunately, these measurements comprise a very inhomogeneous set, since we judged the conventional absorber cone inadequate for the purpose, and since it took until early 1993 to produce a working cold load box. Combined with the periodic changes of the system, this makes a consistent calibration history impossible.

Instead, we rely on the remarkable system stability to use recent measurements to tie the system calibration for the preceding two years. Indeed, while cone measurements are unacceptable because of their uncertain absolute calibration (see section 3.3.2), they are similar from one system calibration episode to the next, and do indeed show a remarkable system stability with system temperatures within 1K over at least one year.

A further complication in understanding previous system calibrations is that the system's non-linearity was not yet known. While this non-linearity has at most a 10% influence on the value of scalable quantities, such as the noise temperature of the CAL, it can have a much more dramatic effect on the estimate of the system temperature, since it involves a subtraction of two large and similar quantities (at least with modern systems at this frequency). For this reason, a careful measurement of the non-linearity is indispensable to a good system calibration and the system temperature measurement—this is described above in section 3.2.12. To include the non-linearity behavior into the system calibration, we simply convert all power readings to 'true' powers using the λ -factor described in the non-linearity section. Similarly, all Cal procedure results are scaled to 'true' CALs in the same way. Here, care must be taken to use the non-linearity law applicable at the attenuation at which the calibration measurements are performed.

The calibration episodes that make up the 'standard' system calibration are data from April 19, 1993 using the TRW amplifier system. Some earlier data are available from March 5,

1993, but a more complete set was taken at the later date. For that calibration, several sets are available with the completed cold load box as well as with the old cone method. Results are very compatible with each other, and do show that cone measurements are consistent with box measurements if the cone temperature is assumed to be 77K. This exonerates the absolute calibration of the cone method, but I prefer the cold load box because of its emission stability, and because full CAL measurements are best conducted using it.

To obtain a similarly good calibration with full characterization of the system's non-linearity for the NRAO amplifier system, the retired but still available NRAO amplifier was swapped back into the receiver in the summer of 1993. This essentially duplicated the NRAO system as it was used during the 1992 observing season, except for the new backend components (*e.g.*, the replaced A/D converter), which are specified to have extremely linear behavior. System calibration measurements for this system were performed on June 24, 1993.

These 'standard' data sets were checked carefully for their consistency with earlier system calibrations. The most consistent system parameters for the NRAO and the TRW systems are given in table 3.5.

TABLE 3.5
CALIBRATION PARAMETERS

Parameter	NRAO System, 1992		TRW System, 1993		Units
	ANT	REF	ANT	REF	
Receiver Noise					
T_{rx}	54.8 ± 1.7	56.4 ± 1.7	33.0 ± 1.6	34.1 ± 1.6	K
Internal					
$T_{rx,i}$	67.0 ± 1.5	69.4 ± 1.5	46.3 ± 1.0	47.1 ± 1.3	K
Horn Noise					
T_{horn}	6.6 ± 2.5	...	5.4 ± 1.8	4.3 ± 2.0	K
Relative Gain	1.000	0.960	1.000	0.964	
Load Isolation					
$\xi_{L,L}$	18.6 ± 0.2	18.4 ± 0.1	18.7 ± 0.1	18.9 ± 0.2	dB
Switched CAL					
$T_{Cal,S}$	0.614 ± 0.003	0.635 ± 0.002	0.545 ± 0.002	0.540 ± 0.003	K

The 'internal' system temperature $T_{rx,i}$ is given to round out the picture; its actual value will depend on the emission temperature of the internal cold loads. By assuming that this emission temperature is well-described by the physical temperature of the internal cold loads as measured by the attached thermometers, the noise temperature of the feeds or horn are calculated. The horn temperature for the REF channel is not available for the NRAO system, since the thermometry malfunctioned during the calibration run.

The values for $T_{Cal,S}$ are taken with Cal procedures of 10s segment length, measuring the switched power level. As explained in section 3.2.13, the value used for the calibration of

the observations must be corrected for the length increase of the CALs during observations, as well as for the Dicke switch isolation. The values of the calibration diodes in the ANT and REF channels are given in table 3.6 with all corrections applied as they are used in

the subsequent analysis. The error bars in that table attempt to cover some of the long-term changes of the CAL, as explained above. The larger error bar in the REF value reflects the apparently larger changes in its value—in both channels the errors from the repeated calibrations of the TRW system have been transferred to the NRAO system as well.

TABLE 3.6
EFFECTIVE V ALUES OF THE CALIBRATION DIODES

	effective T_{Cal}	
	ANT	REF
NRAO System	$0.628 \pm 0.016K$	$0.653 \pm 0.037K$
TRW System	0.557 ± 0.016	0.556 ± 0.037

3.3.4. Efficiencies

Two efficiencies are important for radio antennas (see Kraus 1982 for a discussion of radio astronomy fundamentals): the aperture efficiency and the main beam efficiency. For the observation of discrete radio sources, the aperture efficiency is of overriding interest, since it describes the response of an antenna to a point source; the main beam efficiency, on the other hand, describes the response to an extended source.

For our purpose, we need to know the main beam efficiency to translate observed antenna temperatures into sky temperatures, while we need the aperture efficiency to gauge the effect of discrete point sources on our observed antenna temperature.

Starting from the Nyquist relation (Nyquist 1928)

$$\frac{1}{2} S_0 A_{eff} = kT_A,$$

where S_0 is the flux density of an unresolved radio source that produces an antenna temperature T_A when observed with an antenna of effective aperture area $A_{eff} = \epsilon_A A_p$ that is a fraction of the physical area, we calculate the aperture efficiency as

$$\epsilon_A = \frac{2kT_A}{S_0 A_p}.$$

The main beam efficiency, the fraction of the solid angle of the antenna's sensitivity that resides in the main beam, is then calculated using the main beam's solid angle and the center wavelength

$$\epsilon_B = \epsilon_A \frac{A_p \Omega_M}{\lambda^2}.$$

Efficiencies contain the antenna temperature, which is calibrated using the measurements in the previous section. This gives rise to the fact that the final experimental result is independent of the system's calibration, as long as the observations and the beam efficiency measurements use the same calibration. Since the observed power on a source is related to the observed antenna temperature via the following relation involving the CAL power and temperature

$$T_S = \frac{P_S}{P_C} T_C$$

and since the observed sky temperature of an extended source is related to the true sky temperature via the main beam efficiency

$$T_{Sky, True} = T_{Sky, Obs} / \epsilon_B,$$

we find that

$$T_{Sky, True} = \frac{P_{Sky, Obs}}{P_{C, Obs}} \frac{S_0 P_{C, Src} \lambda^2}{2k P_{Src} \Omega_M},$$

where the CAL temperature is canceled. This implies that no matter how well the system temperature calibration is performed, the accuracy of the final result depends on the accuracy of our knowledge of the beam efficiency and hence of the flux density of the radio source we used in our power measurement of its antenna temperature.

Given the low gain of the 5.5m telescope for point sources, the power measurement is not simple, since none of the available non-variable radio sources of well-known flux and small angular extent generate large antenna temperatures.

The efficiencies of the 5.5m telescope are measured on the basis of the radio source DR21, a luminous HII region with an angular diameter of less than 20"—less than 1/20th of the beam's diameter. Planetary nebulae are advantageous for this purpose, since they emit from a volume large enough to stay non-variable on long time scales, and since they remain luminous at high frequencies because of their optically thin free-free emission, very unlike

the standard extragalactic synchrotron source. The best available value for the flux density of DR21 is given by Klein and Gulkis (1978) in their scale near $\lambda=1\text{cm}$

$$S_{\text{DR21}} = \left[26.70 - 5.62 \log\left(\frac{\nu}{1\text{GHz}}\right) \right] \text{Jy},$$

which gives a flux density of 18.24 ± 0.55 Jy at our center frequency of 32GHz, taking into account the quoted error bar of 3% in the above relation.

Flux measurements of DR21 were performed on various occasions during the observations, but most did not include an experimental accommodation of the system's non-linearity. Therefore, only the data from August 8, 1993 are used below. Applying a correction for atmospheric opacity ($\tau_{\text{atm}} = 0.04$), we find

$$T_{\text{DR21}} = 83.2 \pm 2.8 \text{ mK}$$

for the antenna temperature of DR21 outside the atmosphere (to which the flux density refers), corrected for the system's non-linearity and the Dicke switch isolation. The resulting efficiencies and the telescope gain are given in table 3.7.

TABLE 3.7
TELESCOPE EFFICIENCIES

Aperture Efficiency	ϵ_A	0.531 ± 0.024
Main Beam Efficiency	ϵ_B	0.740 ± 0.039
Telescope Gain	G	$4.56 \pm 0.20 \text{ mK Jy}^{-1}$

4. OBSERVATIONS

4.1. CHOICE OF THE CLUSTER

While the SZ effect has been observed and studied in detail in clusters at medium and large distances, its detection in nearby clusters of galaxies has proved elusive. Some of the reasons for this are given in chapter 2. It is also clear that the measurement of the SZ effect in a nearby cluster holds many significant benefits.

Nearby clusters of galaxies are key objects in our understanding of cosmology, and they are likely to continue to hold much promise. Their advantage is that they are located in a very interesting cosmological region: they do possess ‘cosmological’ distances that are not dominated by the local Hubble flow, but yet they are nearby enough to permit many very detailed investigations of their constituents. Prime clusters for cosmological use are obvious ones: the Virgo cluster, the Coma cluster, Perseus, and clusters in the Perseus-Pisces region.

These objects are easily detectable with modern X-ray telescopes. In fact, Coma has such bright X-rays that it is used as a calibration source for the ASCA X-ray satellite. It is no question that such nearby clusters can be studied with high precision and accuracy in X-rays, a significant factor in the interpretation of SZ observations. In addition, their larger angular size can reveal cluster substructure much more readily than more distant clusters. This applies to X-ray studies, but also to cluster investigations at optical wavelengths, where many studies of cluster substructure have been performed (*e.g.*, Fitchett 1988).

In addition, nearby clusters are likely more evolved than distant ones (*e.g.*, the X-ray temperature structures of 0016+16 at $z=0.55$ and A665 at $z=0.182$; Hughes, private communication). This may apply to their overall parameters as well as to their X-ray atmospheres. Better mixed and therefore more nearly isothermal X-ray atmospheres greatly simplify the interpretation of SZ observations.

Observations of the SZ effect in nearby clusters are therefore often viewed as a prerequisite for a credible cosmological measurement. This applies in the same way to all interpretations of the SZ effect, except for future multi-frequency observations of the peculiar velocity SZ effect.

4.1.1. Telescope Capabilities

As explained in detail in the interpretation section below, most cluster atmospheres can be described by so-called β -laws, in which the radial density profile is described as a function of the central electron density

$$n_e = n_{e0} \left(1 + \frac{r^2}{r_c^2} \right)^{-3\beta/2}$$

where r_c is known as the X-ray core radius. This function is not a random choice; the parameter β provides for a relative assessment of the kinetic energy in galaxies and the gas under the assumption of an isothermal galaxy distribution (see Hughes *et al.* 1988a).

Models with $\beta=0.75$ and a core radius of 300kpc (at $H_0 = 50 \text{ km s}^{-1} \text{ Mpc}^{-1}$ and $q_0 = 0.1$) describe a common bright X-ray cluster. Figure 4.1 shows the response of the two OVRO microwave background telescopes to the peak central SZ effect from such a cluster's atmosphere as a function of its redshift. This response is

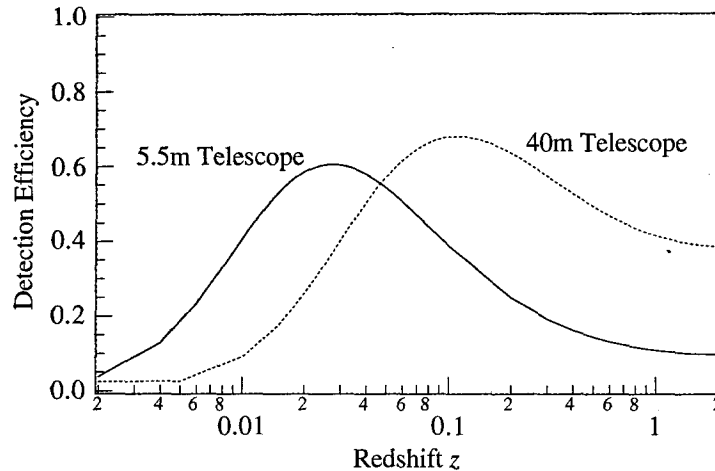


Figure 4.1 Response or detection efficiency of the two OVRO microwave background telescopes to the peak central SZ effect in a typical cluster (see text) at various redshifts.

not equal to one or even constant, since the telescope's position switching cannot place the reference beams entirely outside the cluster, while the central peak effect is diluted by the central beam width. This figure was calculated with an estimated beam size of 7.5 arcmin for the 5.5m telescope and a width of 1.77 arcmin for the 40m telescope; both numbers are slightly out of date, but the results remain substantially unchanged.

4.1.2. Requirements

This experiment was not intended to survey a complete sample of clusters initially, but because of its difficulty merely aimed for a measurement of the effect in one good candidate. The choice of this cluster is dictated by the following considerations. Since the SZ effect is proportional to the integral of the product of temperature and density along the line of sight, we should choose to observe a cluster with a hot, dense, and large X-ray atmosphere.

A further requirement is a high level of regularity in the cluster's X-ray structure. This immediately excludes such distorted clusters as Virgo. If the cluster possesses a strong cooling flow, the interpretation of the data becomes substantially more difficult. Not only is the X-ray surface brightness polluted by lower-temperature gas that contributes little to the SZ effect (this effect is similar to clumping), but a cooling flow could produce sufficient amounts of free-free emitting gas at intermediate temperatures (see Tarter 1978) to corrupt the measurement of the SZ effect.

On a practical observational level, the target cluster must be free of strong radio sources. This does not mean that a cluster must be completely clean, but the presence of 3C84 in the middle of the Perseus cluster (A426), for example, guarantees that this cluster will not be observed with microwave background methods. Unfortunately, many nearby clusters have a high level of radio source contamination, and many high-temperature clusters were rejected for observations because of this reason.

An obvious cluster that fulfills these requirements is the Coma cluster, A1656. Its redshift of 0.0235 (Sarazin *et al.* 1982) places it directly at the peak of the 5.5m telescope's sensitivity curve, as shown in figure 4.1. Indeed, it has been the prize of SZ observations from the start, and has been the not-so-secret candidate Sunyaev and Zeldovich had in mind when they proposed observations of their effect.

The cluster has a comparatively regular structure, being cD-dominated with Bautz-Morgan Type II and Abell richness class 2 (Abell *et al.* 1989). Recent pointed Rosat observations (White *et al.* 1993) do show that some structure is discernible in the cluster's surface brightness images, but this structure is minor in comparison with most other clusters. Of course, the interpretation of SZ observations will have to take such structure into account, but it is not pronounced enough to prevent useful observations and interpretations to be made at this stage.

Coma is hot and dense: the most sophisticated hybrid X-ray model to date (Hughes *et al.* 1988b and Briel *et al.* 1992) shows that it has a central (isothermal) temperature of 9.1keV and a central density of about $n_0 = 3 \cdot 10^{-3} \text{ cm}^{-3}$. The temperature is quite respectable; very few known clusters reach such energies. While the hottest known X-ray cluster, A2163, has a temperature of $14 \pm 1 \text{ keV}$ (Arnaud *et al.* 1992), such high temperatures are unusual, and even temperatures as high as 8keV occur only in rich clusters. The X-ray model of the cluster is discussed in more detail below.

The radio environment of the Coma cluster is reasonably well-known, and it contains no very strong radio sources. Most of the brighter radio sources in the cluster are associated with its member galaxies (Venturi *et al.* 1990), and possess steep radio spectra. Only one bright flat-spectrum source is known, 5C4.105, which sports a flux density of 210mJy at 18.5GHz (see section 6.2). It is also the third-brightest radio source in the cluster at 1.4GHz; the brightest one, 5C4.81, has a flux density of about 370mJy at that frequency (Willson 1970), but decreases with a spectral index of more than -1 toward higher frequencies. Fortunately, the flat-spectrum source, whose spectrum above 18.5GHz is unknown, lies well apart from the observing patterns of this experiment (36' from position C and 10.5' from position N; see section 4.3 for their definition).

The center of the cluster contains a further radio source: an extended radio halo with a diameter of about 20', known as Coma C. While the luminosity of this radio halo is enormous at low frequencies, it possesses a very steep spectrum ($\alpha = -1.3$; Venturi *et al.* 1990), which lowers its effect on our observations to a few microkelvins (see section 6.3).

But one of the real advantages of observing the Coma cluster is that it is well-known at all wavelengths. Not only are we well-acquainted with the cluster itself—its context is also well-studied, since it lies at the heart of the CfA redshift survey. Models of the X-ray atmosphere go beyond the simple isothermal sphere with a beta-model density law; the hybrid model developed by Hughes *et al.* (1988b) provides a much better fit not only to central data, but also to the cluster's X-ray atmosphere at radii of 40' and even up to 100' (Briel *et al.* 1992).

4.2. OBSERVING TECHNIQUE

As was demonstrated in the instrumentation chapter, great care was taken to ensure that the telescope enjoys maximum freedom from systematic effects. Still, our desired maximum level of systematic effects of $10\mu\text{K}$ places extremely stringent demands on instrumental performance. Since it is virtually impossible to ascertain that all systematic offsets and biases have been eliminated at that level from the system, differencing schemes are usually employed during observations. In typical single-dish radio observations, two levels of differencing are employed. We felt that the added observing time and the reduced sensitivity of an additional level of differencing was well worth the added insurance. This is especially true in light of the single largest source of systematic error in our system: differential ground pickup. Measured to have a peak value of $150\mu\text{K}$ at some zenith angles, this bias varies as a source is tracked over the sky. Clearly, the additional level of differencing should be designed to cancel this remaining differential ground pickup specifically.

Differencing

In our observations we use a three-stage differencing technique. The first two stages are known as double switching, while the third is aimed at controlling residual differential ground pickup.

The first level is achieved by constant Dicke switching between the two beams on the sky at 500Hz. This eliminates all constant offsets (such as total power) and reduces linear drifts common between the two sides to the amount incurred during one Dicke switching period. In the case of non-common drifts, the system's sensitivity to these is reduced to the sum of the drifts during one Dicke switching period. Even non-linear fluctuations are reduced as long as their time scale is longer than that of the Dicke switching. Therefore, faster switching times yield lower effective offsets and are highly desirable. This is illustrated in section 3.2.6.

As explained above, Dicke switching switches the input to the first-stage amplifier between the receiver's two feeds. This means that all elements behind the Dicke switch are identical, and that Dicke switching will reduce drifts in those elements very well. On the other hand, the two RF sides from the Dicke switch to the front of the feeds are independent—the Dicke switch will not be able to reduce fluctuations in those elements as effectively.

The parts of the system before that are a more complicated question. Certainly, the two beams share the same antenna, but the illumination patterns overlap imperfectly. The situation is similar with the column of air in front of the telescopes: the closer sections of the beam overlap much more than the farther ones. In fact, the two beams of the telescope diverge completely at a distance of 1200m; from then on they traverse completely independent columns of air. This distance is not very large and explains the strong weather sensitivity of the 5.5m telescope.

Still, Dicke switching is probably the most effective level of differencing in our observations. Among the biases that are reduced very well by Dicke switching are short-term receiver gain fluctuations, 60Hz interference, and the bulk of ground pickup and atmospheric effects.

The second level of subtraction is accomplished by position switching in which the sky source is alternately placed in the main and the reference beams in a symmetric pattern: REF-ANT-ANT-REF. The antenna is physically moved by one beam separation between individual segments. This procedure, described as the Flux procedure in section 3.2.1, cancels all remaining linear drifts in the apparent brightness temperature, leaving us sensitive only to temporal and spatial curvatures. This technique implies that our effective sensitivity on the sky is distributed over three points: the main beam as well as the average of two reference beams offset from it in azimuth. The consequence of this azimuthal separation is that while tracking a source over a range of hour angles, the reference beams will spread over parallactic angle to form reference arcs flanking the central field. This parallactic angle ψ is defined as

$$\tan \psi = \frac{\cos \varphi \sin H}{\sin \varphi \cos \delta - \cos \varphi \sin \delta \cos H},$$

where φ denotes the telescope's geographic latitude and δ and H are the object's declination and hour angles, respectively. In the analysis below, the parallactic angle is restricted to $\pm 90^\circ$ because of the effective beam's azimuthal symmetry, but this redundancy could in principle be retained, especially for non-symmetric beams. With this definition of the parallactic angle, a source with $\delta < \varphi$ will rise at negative parallactic angles, transit at 0° , and set with a positive angle. In the projection of the beam pattern on the sky, a zero parallactic angle means that the beam pattern is parallel to right ascension, while the parallactic angle's direction corresponds to its mathematical usage.

The third level of subtraction attempts to remove residual differential ground pickup, which appears as a real signal, but one fixed with respect to the ground rather than the sky. For this procedure, we alternate between the target field and a blank reference field 15 minutes away in time, so that the telescope tracks the two fields over precisely the same paths relative to the ground. For this procedure, both leading as well as trailing reference fields can be used. Below, I will refer to a 15-minute track on the target field as an ON scan, while its colleague is known as an OFF scan, regardless of whether it leads or trails the target field.

In the case of the observations described here, two basic patterns were employed: MAIN/TRAIL; and LEAD/MAIN; each in two versions shifted in hour angle to even the parallactic angle coverage. It is also possible to use a single LEAD/MAIN/TRAIL pattern, shifted into three versions to equalize parallactic angle coverage. Such a scheme is employed in SZ observations on the OVRO 40m telescope.

Each segment had a length of exactly 15 minutes and contained between 4 and 6 individual double-switched flux measurements, depending on the zenith angle. The segments are synchronized such that each flux measurement in an ON scan covered precisely the same azimuth and zenith angle range as the corresponding one in an OFF scan.

As a result of this differencing procedure, observations of a position in the cluster are sensitive to 9 regions in the sky.

Observing Schedules

All observations were controlled by schedules that are run by the telescope's control software. These files are an exceedingly convenient way of observing. Not only can they run unattended for months (in principle), but they are also the only method of assuring consistent timing and minimal observer error. In addition, they preserve an observers' sanity. An example of one of the schedule files used during the 1993 season is given in appendix C.

For the third level of differencing, care must be taken that the observing procedure allows sufficient overlap between corresponding flux measurements in the ON and OFF scans. At times, the wind is high so that the telescope does not acquire the source with the same speed it usually has, or there are encoder or drive troubles that delay acquisition and tracking. A good observing schedule should be robust under these circumstances. In fact,

the standard SZ observing schedule (see appendix C) ensures that no scan will be allowed to start if it is delayed by more than 30 seconds. If such a delay should occur, the scan will simply be skipped. This does have negative consequences. For instance, it means that the scan corresponding to the skipped one in the ON/OFF differencing will be discarded in the later analysis, since it cannot form a pair with the missing scan. However, this tradeoff is far preferable to the possibility that the timing of all subsequent scans is corrupted.

Each scan contains a standard sequence of radiometry procedures to ensure consistent data as well as guaranteeing that all relevant data are recorded. First come the Flux measurements with varying parameters, as explained below. For synchronization reasons, they come before all other procedures. They are followed by one Cal procedure to determine the current system gain by measuring the apparent power of the ANT calibration noise diode. During this procedure, the load switches are pointed to the internal cold loads so that the calibration value is measured on top of a very stable power level. After the Cal procedure terminates, the switches are returned to their usual observing position.

This is followed by two average procedures measuring total power both on the sky and on the internal cold loads to monitor the system temperature level and to correct for the system's non-linearity. For the most part, the sky total power value measures the current atmospheric emission, a useful measurement during analysis. A measurement of the zero level completes the procedures of a scan.

The timing of these procedures is 89s for the Cal procedure (20s per segment, 3s idle) and 8s for each average and zero procedure.

Flux Parameter Optimization

As a source is tracked across the sky, the length of a scan stays constant. However, at different zenith angles the telescope can take very different times to accomplish the position switching between flux segments as well as the ground spillover switching to the reference fields. Especially in a cluster like Coma that transits at a zenith angle of 9.5° at OVRO, the slew times at transit are much longer than at more moderate zenith angles. Because of the limited time available per segment and because of the desire to control the timing of the flux measurements exactly to ensure proper synchronization between corresponding segments, the timing parameters for the flux measurements should be given explicitly in the observations. Since it would be wasteful to retain the long slew times required at transit for

moderate zenith angles, the timing parameters were adjusted as a function of zenith angle. The 1992 season saw a first version of this dynamic timing, while a more sophisticated version was used in subsequent observations.

Tests to determine the time required for the position switching within a Flux measurement as a function of the scan's minimum zenith angle Z_{min} found that this idle time parameter (see radiometry procedures in 3.2.1) τ_i is

$$\tau_i = 6s + \frac{3.5s}{\sin Z_{min}}$$

Using a fixed allowance of 113s for the other radiometry procedures, a slewing allowance to get from the ON to the OFF field and vice versa of 60s at $Z < 20^\circ$ and 45s otherwise, and a 'slop' of 15s, this variability of the idle time can be used to find a nearly optimal choice of parameters for the remaining Flux procedures. The parameters listed in table 4.1 were calculated to maximize

the observing efficiency ϵ , the fraction of time spent on Flux integrations during a full scan. To avoid concentrating on long segment integrations S , which of course maximize efficiency because of the fewer slew times involved, an upper limit of 25s per segment was used, except in the case of transit observations. The initial table was then fine-tuned after a

TABLE 4.1
OPTIMAL FLUX PARAMETERS, 1993

Scan	H	ZA Range		FLUX Parameters			ϵ
		min	max	S	I	R	
1	-3 30			23s	12s	6	61%
2	-3 15			23	12	6	61
3	-3 00			23	12	6	61
4	-2 45			22	13	6	59
5	-2 30	30.6°	32.2°	22	13	6	59
6	-2 15	27.3	29.2	22	14	6	59
7	-2 00	24.4	26.3	21	15	6	56
8	-1 45	21.6	23.4	21	16	6	56
9	-1 30	18.8	20.6	26	17	5	58
10	-1 15	16.2	17.8	24	19	5	53
11	-1 00	13.8	15.2	23	21	5	51
12	-0 45	11.6	12.7	21	24	5	47
13	-0 30	10.1	10.8	28	28	4	50
14	-0 15	9.4	9.6	27	30	4	48
15	0 00	9.4	9.7	27	30	4	48
16	0 15	10.1	10.9	27	27	4	48
17	0 30	11.7	12.8	21	24	5	47
18	0 45	13.8	15.2	23	21	5	51
19	1 00	16.1	17.8	24	19	5	53
20	1 15	18.8	20.6	26	17	5	58
21	1 30	21.6	23.4	21	16	6	56
22	1 45	24.4	26.3	21	15	6	56
23	2 00	27.3	29.2	22	14	6	59
24	2 15	30.3	32.2	22	13	6	59
25	2 30	33.2	35.2	22	13	6	59
26	2 45	36.2	38.1	23	12	6	61
27	3 00	39.2	41.1	23	12	6	61
28	3 15			23	12	6	61
29	3 30			23	12	6	61
30	3 45			23	12	6	61

1993 season scan scheduling parameters. H is the hour angle at the start of the scan, which has R Flux measurements with S integrations per segment and a slewing allowance of I . The overall efficiency ϵ gives the scan's fraction of pure integration time.

few days of observations. It can be seen that these parameters maintain above 50% efficiency in all scans save the three transit ones.

Other Considerations

In the 1992 season, the hour angle range of -2.5 to 3.5 hours was used for observations; in the following season, this range was expanded to -3.5 to 4 hours. In both seasons, an attempt was made to even out the proportions of data with the leading and with the trailing reference fields, along with smoothing out the data's hour angle distributions. When an OFF field is observed, the hour angle coverage is identical to that of the ON field, while a 15 minute segment of the entire hour angle coverage is being missed. For this purpose, two schedule versions are needed for each reference field, shifted to each other by those 15 minutes. Therefore, a total of 4 schedule files are used to accomplish the observations. Appendix C shows an example of the 'Trailing OFF Field, Odd scan number' schedule used during the 1993 season.

The schedule also sets up and maintains the multi-channel chart recorder, so that the observing conditions are constantly monitored during all observations. When SZ observations are not running (*e.g.*, Coma is out of its observing hour angle), the telescope is used for observations of intrinsic microwave background fluctuations. Since the chart recorder is maintained during those observations at the same settings, a continuous record of observing conditions is available.

Four channels are monitored on the chart recorder: both total power channels, ANT and REF, and the switched power are the radiometry channels. The fourth channel is the temperature of the digital backend to give an indication of the receiver electronics' temperature stabilization.

The purpose of the chart record was not data keeping, since that is accomplished very effectively by the control system. Rather, it can be used during observations as a quick way to gain an overview of current observing conditions, and to spot problems, such as equipment failure and software problems easily. The chart records are also used as a retrospective record of the observing weather, and are used in the data analysis.

Observing

Observations are carried out continuously during the observing season, which stretches from the beginning of February to the middle of May for the Coma cluster. Because of this duration, it is impossible for one person to attend observations all the time. Nor is it necessary. We have tried to keep at least one member of the microwave background group in residence at the observatory during the entire microwave background season, to keep track of the systems and to spot problems as soon as possible. During the 1993 season, I spent between one-third and one-half of my time at the observatory.

Observations are also usually carried out under all weather conditions, unless periods of bad weather permit other test data to be taken. The removal of data taken during bad weather is accomplished later. This ensures consistency of our weather editing procedures, rather than relying on the observer's subjective feeling at the time. Of course, plenty of additional data are recorded to be able to judge the weather performance later. With the 5.5m telescope weather conditions can be judged mostly by the data quality alone. However, several strands of information such as WVR data, chart records, weather logs, and log file entries by the observers help to form a much more consistent picture.

4.3. OBSERVING POSITIONS

Once a target cluster was selected, the question was where to point the telescope. With single dishes only one value can be derived from each pointing position. The ideal case would be to obtain a full grid of SZ observations over the cluster; while this remains the ultimate goal, first observations must concentrate on fewer pointings, since much observing time is needed to detect the effect with sufficient solidity. Two positions could be observed in the first season, and a third was added during the following winter.

Clearly, a position close to or at the center of the cluster has to be observed. I chose to observe a position offset from the cluster center by $2' 07''$ to the south and by $0^m 11^s$ to the east of the cluster's published X-ray center (Abramopoulos *et al.* 1981), giving a total offset radius of $3' 22''$. This was done to avoid pointing in the immediate vicinity of the cluster's central radio sources without moving away appreciably from the cluster's central

SZ peak. This offset from the cluster center must be taken into account in the interpretation of the observations below.

TABLE 4.2
OBSERVING COORDINATES

Name	Position	Main Field		Leading Reference		Trailing Reference	
		α	δ	α	δ	α	δ
Center	C	12 57 30.0	+28 11 00	12 42 28.1	+28 11 12	13 12 31.9	+28 10 44
North	N	12 57 19.0	+28 43 07	12 42 16.9	+28 43 19	13 12 21.0	+28 42 51
Extension	CS	12 54 31.0	+27 43 40	12 39 29.1	+27 43 52	13 09 32.9	+27 43 24
X-ray Center		12 57 19.1	+28 13 07
Abell Center		12 57 24.0	+28 15 00

NOTES Coordinates of the three SZ positions observed with the 5.5m telescope, along with the two types of reference OFF positions. Also given are the nominal X-ray center (Abramopoulos *et al.* 1981) and the Abell center (Abell *et al.* 1989). All coordinates have epoch 1950.

The first season also saw the start of observations to compile a north-south profile of the effect. Since there was insufficient time to compile a full profile, I concentrated on the northernmost field to observe a point in which the expected SZ effect is very small. This northern point lies 30' from the cluster's X-ray center—it is not in line with the central pointing, which should not introduce further complexities into the interpretation.

With successful observations in the first observing season and considerable speculation about possible parallactic angle structure of the effect, a further position was selected for observations: the center of the south-western expansion of the cluster's atmosphere, about 1° away from the main cluster center. This is also a region of increased radio source density (*e.g.*, Venturi 1985) and harbors the large, extended radio source 1253+275 at its center. I estimated its center position from the maps published in White *et al.* (1993).

4.4. DATA

Observations of the SZ effect were made during the three winters from early 1991 to the spring of 1993. Because of the discovery of the ground pickup problem described in chapter 3, all observations from the first season had to be discarded. Unfortunately, this season had the best observing weather in the three years of observations.

The second season in the winter and spring of 1992 provided the detection of the effect and remains our best data set. The third season, with a more sensitive system, was affected by

very unfavorable weather conditions, but does contain some high-quality data. It provided the confirmation of the detection and some additional data.

With the TRW system, including full ground subtraction, we can expect to detect the SZ effect at the 3σ level in a cluster like Coma in 9–10 hours of telescope time, a little more than one night of observing. This figure takes into account the rise of the dynamic system temperature by a factor of 2.5 on long time scales, an average observing efficiency of 56% (see section 4.2), and a target standard error of $50\mu\text{K}$. Of course, this requires excellent weather conditions.

Observations were made for substantially longer periods, so that several positions in the cluster could be measured to high accuracy.

1992 Season

Observations of the SZ effect in Coma began on February 14th, 1992 and ran through May 4th, 1992. Starting with a fully ground-subtracting observing schedule with dynamic parameter selection, the center position was observed for 6 hours centered on transit. Observations of the center position continued to obtain a firmer result and to improve the statistics. Finally, on April 20, 1992, I switched observations to the northern position N.

During late winter and early spring, Coma is overhead at night. This determined the choice of the observing period, since the data are substantially more stable at night than during the day, especially in marginal conditions.

This season coincided with an abatement of California's seven-year drought. While it had a mixture of good and bad weather, it contained no long periods of the cold, pristine weather conditions that have always been the backbone of ground-based microwave background observations. Atmospheric observing conditions deteriorated markedly toward the spring, with rising temperatures and a continuous storm track. For this reason, combined with the shorter total observing time on position N, the result from that field has much lower quality than that from position C.

1993 Season

During this season, which ran from February 15, 1993 to April 18, 1993, the main objective was to confirm the detection of the SZ effect in position C which was obtained during

the previous season, and then to extend observations to other points in the cluster. After this was accomplished, I switched to a 'hybrid' scheme in which position N was observed during the same night as the center of the cluster's extension, position CS. Observations with this scheme started on March 3 and continued to the end of the season.

This season saw substantially worse weather than before, with California's drought finally ending. Only very rarely did cold conditions coincide with completely clear skies—for the most part, the sky was covered with thin cirrus and the valley's air was filled with haze. Again, atmospheric conditions, and therefore data quality, deteriorated noticeably during the season. In fact, December provided some excellent observing conditions, which could not yet be used for these SZ observations.

Even though this season failed to provide the expected increase of observed cluster positions, it did succeed in the vital goal of confirming the observations of the previous year.

5. ANALYSIS

5.1. OVERVIEW

Conceptually, the data reduction occurs in two steps. First, the data are edited with care to avoid the introduction of biases into the final result. In this experiment, this has been ensured by restricting all editing steps to the use of standard errors only. Methods that use the data values themselves are possible, and have been used successfully in other microwave background work (*e.g.*, Readhead *et al.* 1988; Myers *et al.* 1993). In this experiment, an additional level of differencing had to be accomplished after the editing stage, and I decided not to apply those methods.

The parameters of the editing algorithms were determined with a test data set representing a cross-section of observing conditions. The parameters derived from that step were applied consistently and uniformly to the entire data set.

After the editing stage, the data are summarized and combined to form the result. It is at this stage that the statistical reliability of the data set can be explored.

A brief description of the analysis package appears in appendix D.

Definitions

The actual data consist of procedure results from the Flux procedure. These are called *data points* and are the fundamental data unit provided by the telescope. Each data point taken while pointing to the ON field has a corresponding data point taken while pointing to one of the reference OFF fields; the two form a *data pair* and constitute the fundamental data unit used in the calculation of the result. A group of data points that was taken contiguously while pointing at the same field is called a *half scan*, which typically has a length of 15 minutes and contains between 4 and 6 data points. From the standpoint of the calibration, the half scan is the fundamental data unit: each half scan contains one Cal procedure to determine the current system gain, as well as other radiometry information. The calibration

information is interpolated over the half scan to determine its value at the time of each data point.

Two corresponding half scans, one in the ON position and the other in the OFF position, are referred to as a full *scan*, which therefore consists of a series of data pairs. Depending on the observing configuration, the ON half scan may occur before or after the OFF half scan. A full night's observing, then, consists of a number of full scans (half scans will be discarded if not matched to become a full scan), and is called a *day*. A number of days, hopefully large, makes up a season, which is the largest data sub-unit before the final result.

Appendix E contains a summary of the statistical definitions used in the analysis.

5.2. DATA EDITING

Editing of the data is perhaps the most critical aspect of the analysis of microwave background observations. While careless editing can introduce severe and subtle biases into the results, some editing is necessary because most sources of bad data produce non-Gaussian values that would corrupt the result if treated in the usual Gaussian fashion. Therefore, editing methods should identify all bad data without biasing the data by selective removal.

Several principles have been applied to obtain such methods. First, editing methods attack specific causes of corrupt data. Further, no editing procedure is allowed to operate on the actual flux values used in the final result. Instead, they must use quantities that will reliably indicate bad data while not using the flux values themselves. As indicated above, they use the internal error estimates of varying lengths of data.

In addition, all editing procedures are arranged to preserve complete ON/OFF data pairs to ensure proper ground subtraction. A final requirement of editing methods is robustness. Small changes in the editing parameters should not affect the result strongly.

Before the main editing algorithms are allowed to operate, data with obvious problems are removed. This includes a perusal of the computer logs that were entered by the observer and a look at the chart records for clear signs of trouble such as system warm-ups and scheduler problems. In addition, all saturated data are removed, since they can simulate exceedingly good weather by their internal error of 0 counts.

Another level of first-stage editing was performed implicitly. Data taken entirely during the worst weather conditions were not even included in the editing process. This was done with the certain knowledge that the automatic editing procedures would reject these data completely. If there was any question about the usability of a day, or if some part of it appeared usable, the data were included in the editing.

5.2.1. CAL Editing

Before any calibration is applied to the data, measurements of the CAL can be edited. This is the only case in which manual graphics-based editing has been permitted in the editing routines. Typically, the CAL values from a day of observing are plotted and individual measurements can be removed. The purpose of this editing step is to ensure that an occasional glitch in the CAL value will not corrupt legal data. From my experience with the telescope, the value of the CAL is extremely stable and varies only within a small envelope over several hours, reflecting slow, diurnal changes in the system gain. Sharply discrepant individual values of the CAL are therefore never indicative of a true change in the system gain, and it is quite necessary to remove them. Mostly, such discrepant values are caused by bad weather conditions, which are 'visible' even through the load switches during the Cal procedure.

This manual editing method for the CALs is necessary, since no automatic method has been devised. Given the very slow change of the value of the CAL and the fact that missing CALs affect the data only by scaling with a slightly different interpolation, such a method was not found essential.

5.2.2. Overlap

With the ground-subtracting observing technique described above, we are counting on measuring the same ground pickup when we point at the same position with respect to the ground at very similar times. Conversely, we cannot count on subtracting ground pickup properly if we do not point to exactly the same position with respect to the ground. In other words, the altaz position of an ON data point must overlap that of its corresponding OFF data point.

Data pairs fail to overlap when the observing schedule loses synchronization. This can be caused by drive problems or by wind loading (this is not as severe on the 5.5m as on the 40m telescope).

It is hard to judge how precise the overlap must be to cancel ground pickup satisfactorily. Certainly, the sidelobes through which ground radiation enters the system are more diffuse than the main beam, which is the telescope's most pointed lobe—only in that direction do all reflections over a large lateral span add in phase. It therefore follows that overlap of a fraction of the main beam's size must be sufficient for proper ground pickup removal. Because of the resolution of the azimuth and zenith angle records in the log file of 0.01° , this overlap criterion should not be smaller than $2'$, which is the parameter chosen for the analysis. Such a stringent criterion, compared to the beam size of about $7'$, runs the risk of being overly conservative—in practice, mostly data pairs with clearly unacceptable overlap are removed by this procedure.

5.2.3. Bad Weather Removal

Since observations take place essentially at all times, some way must be found to remove data taken during unfavorable atmospheric conditions. This after-the-fact removal is unavoidable, since the weather conditions are usually best judged with a broader and more consistent view of the data and of concurrent weather records. As indicated above, bad weather conditions do not only increase the thermal noise of the observations (by increasing total atmospheric emission), but they introduce strongly non-thermal fluctuations into the data that can bias the result and impact the averaging of good data. The latter is far more serious.

How, then, do we identify the influence of bad weather conditions? Apart from plain weather logs, whose importance is undisputed, but which cannot be kept at night, there are two basic approaches to this question. The first uses the fact that with increasing thermal emission from the atmosphere, the associated non-thermal fluctuations grow as well. By excising data which were taken during periods of atmospheric emission higher than some threshold, data from bad weather can be removed. This approach was taken very successfully with microwave background data in the OVRO 40m experiments (Readhead *et al.*

1988, Myers *et al.* 1993). Because of its larger switching angle, the 5.5m telescope is more weather sensitive than the 40m telescope, so that this method could not be used.

The approach used is to measure the direct non-thermal impact that bad weather has on the data themselves. It is a two-step method that first excises all data with increased dynamic system temperature. The second step recognizes that the weather can start biasing the data before its non-thermal effects become directly visible in the data—this results in variable margins of seemingly good data to be removed around groups of rejected data.

Thermal Cutoffs

The first step removes all data points whose dynamic system temperature is increased over the nominal system temperature by more than a factor of 2 on timescales of a flux measurement (typically 100s). The dynamic system temperature takes into account the zenith angle dependence of the atmospheric emission. For this, the air mass is simply calculated as

$$\chi_{atm} = \sec Z.$$

The calculation also uses the actual integration time expended for each data point, which changes dynamically from scan to scan.

This editing technique works extremely effectively and removes all sections of the data with noticeably elevated noise levels. As will be seen below, the vast bulk of rejections in the data was accomplished by this criterion; the margin criterion described below added only a modest number of rejections.

Group Editing

Still, observing experience has shown that periods of noticeably bad weather were usually preceded by periods during which the cloud cover slowly increased, without yet producing a significant increase in the short-term fluctuation estimate. However, there is little question that such periods could still affect and possibly bias the data.

Based on the apparent pattern that weather conditions usually took a few hours to deteriorate and that short periods of bad data were usually only caused by a relatively local disturbance, the second part of the bad weather removal employs a simple-minded approach. All

contiguous groups of data that were rejected on the first pass by the thermal cutoff criterion were recognized and surrounded during a second pass with margins of rejected data on both sides, each of which equaled the length of the thermally rejected group. As intended, this introduces some measure of dependency on the duration of the disturbance. The maximum length of such a margin was set to 2 hours, roughly the length of bad premonitory weather, to ensure that not all data from a night are lost if the weather cleared up quickly at some time. No rejection margins were formed around isolated single data points cut by the thermal criterion.

This second weather editing stage is mainly intended as a conservative safety measure; its use is justified given the absence of an editing procedure that measures long-term fluctuations (see below). In addition, it was intended as a group-finding procedure to eliminate short sections of seemingly low-noise data in the midst of long high-noise regions. The margin criterion removes those sections very successfully.

This safety method did not result in vast amounts of rejected data—quite the contrary, it operated mostly within regions of significantly bad data. In the 1993 season, for example, 1571 data points out of a total of 4276 were rejected by thermal editing procedures. The number of points that were actually rejected because of the margin criterion alone was 398, but this fraction is reduced even further once complete data pairs are considered: out of a total of 2223 data pairs in the season, of which 1036 were rejected by thermal editing, only 68 data points were rejected because of the margin criterion alone. A good safety to have with low cost to the data!

5.2.4. Pair Completion

With a fully ground-subtracting observing technique, some method must be found to preserve proper cancellation. If the cancellation was accomplished merely by subtracting the average OFF value from the average ON value, we would have to be sure that all sections of the hour angle track (along which ground pickup will vary) are subtracted with equal weight. Even if this could be accomplished by subtracting the average hour angle profiles from each other, data that are supposed to cancel each other would still be taken during different weather conditions, in which the magnitude of ground pickup will vary due to the change in physical temperature as well as changes in ground emissivity. This

means that biases could still creep into the analysis. For this reason, exact ON/OFF data pair completion is enforced at all times.

5.2.5. Editing Methods Not Used

Perhaps a description of the editing methods used in the analysis is best completed by a discussion of possible editing methods that were not employed.

As explained above, no editing was performed that made decisions based on the flux values; instead, only the scatter inherent within one data point was used. Further, it goes without saying that at no stage were data removed by graphical editing or otherwise manually, unless there was a clear reason to do so. Only 6 data points out of a total of 12475 were removed manually because of unusual switched power behavior visible in the chart record.

There are some editing methods, described below, that could be applied to the data in this experiment with the appropriate caution. Their main drawback is that they suppress structure in the data. Such structure may be perfectly legitimate: differential ground spillover that is subtracted later by the pairing operation; the influence of discrete radio sources; and structure in the SZ effect itself, which would be seen in the measurements' parallactic angle distribution. However, such a leveling effect is most likely small, since the size of the expected structure is much smaller than the statistical distribution's width, and it might be instructive to apply such methods to the data in this experiment. This is especially true since the antenna temperatures of discrete radio sources are very small with the 5.5m telescope, and would fall below the width of the distribution for any cluster of galaxies that has small enough radio sources to be considered for observation. With ever decreasing system temperatures, however, the width of the distribution continues to decrease, which may invalidate such methods in the future.

These methods fall into two groups: the first is to remove data values that are outliers in the distribution; the second involves measuring the noise in the data on time scales longer than the 100s of a single flux measurement.

Outlier removal techniques recognize the fact that the actual distribution of the data is not Gaussian, but that it has extended tails of invalid data. Often such outliers possess small internal error bars, which makes them very difficult to eliminate with more conservative

editing methods. To remove outliers in a way that does not bias the mean of the data, iterative techniques are used that decide whether a point is an outlier based on a mean that is recalculated after data are removed, so that a bias cannot be frozen into the data simply by where the initial mean may have been. A statistical correction is then applied to the resulting error bars, which will be artificially depressed by the removal of the data's Gaussian tails.

The second group of methods recognizes bad data by determining the fluctuations of the data values (not their initial individual noises only) in a window of the desired length. Data are then cut on the basis of this window. This method has not been used in the present experiment, since the method's null hypothesis is a signal that is constant with hour angle, zenith angle, or parallactic angle. However, changes of the data point values against those quantities are entirely expected, from the subtraction of differential ground spillover to the shape of the expected SZ effect profile with parallactic angle due to structure in the SZ and due to discrete radio sources. If such a scheme were applied to complete data pairs to minimize its effect on the differential ground spillover subtraction, the time scale is no longer clearly defined. Still, a more sophisticated version of such an algorithm might be workable and would help to identify further sections of bad data.

5.2.6. Rejection Statistics

The effect of all the above editing methods in reducing the number of data points and complete data pairs is shown in table 5.1. It is apparent that the bulk of the rejections was accomplished by the thermal cutoff criterion; the other methods simply act to avoid occasional errors in otherwise excellent data. The large number of thermal margin rejections describes mostly data that were already rejected by the thermal cutoff criterion, or had pair partners that were.

TABLE 5.1
NUMBER OF REJECTED DATA PAIRS

Criterion	1992		1993	
	C	N	C	N,CS
Manual Rejection	6	0	0	0
Data Dropout	30	0	68	0
Missing ON Field	180	26	87	0
Missing OFF Field	78	27	83	0
Altaz Overlap	161	36	67	28
Thermal Cutoff	502	213	1196	439
Thermal Margin	370	119	887	367

Table 5.2 shows the total number of data points and data pairs that were taken in each season in each position. Since some days with visibly and evidently bad weather were not even included in this analysis in the 1992 season, the fraction of good to total data is not very significant. Still, it is apparent that the weather during the 1993 season did not allow many data points to remain in the edited data set.

TABLE 5.2
TOTAL AND USED DATA POINTS AND PAIRS

Position	Data Points		Data Pairs	
	total	good	total	good
1992 C	4354	3489	2305	1518
N	1743	1352	898	583
1993 C	4276	2314	2223	955
N	1080	...	540	295
CS	1022	...	511	291

NOTE the number of good data points is not available separately for the 1993 N and CS positions.

5.2.7. Robustness

Does the editing and the choice of editing parameters introduce any bias in the final result? Since the editing methods were designed with specific targets in mind, it is expected that such a bias does not occur. However, especially considering the choice of thermal editing parameters, a certain degree of independence of the result from the editing parameters must be shown.

Fortunately, all data sets fulfill this requirement to great satisfaction. Figure 5.1 shows the averages of the 1992 position C and all 1993 position data sets as a function of the thermal cutoff parameter. It is clear that the result is very insensitive to changes in this parameter; the departure at small values is due to the fact that the majority of data is rejected at such a stringent criterion. However, the cutoff value of 2 that was actually used in the analysis lies within remarkably stable regions.




Figure 5.1 Sensitivity of the weighted results to the thermal cutoff parameter. Shown are the 1992 season data on position C, as well as the 1993 season data on all positions.

5.3. RESULTS

5.3.1. Correction Factors

Several correction factors must be applied to the values that are produced by the telescope. Some of these are contained in the calibration and are therefore applied implicitly, while other corrections must still be made explicitly. The factors inherent in the CAL value include the effect of non-linearity on the system calibration as well as the absence of ideal Dicke switch isolation. In addition, the effect of atmospheric opacity has already been corrected within the analysis program (using $\tau = 0.04$ for the retained good-weather sections), since it is zenith angle dependent and therefore varies from one data pair to another.

Further corrections to convert instrumental units to an equivalent SZ temperature difference at zero frequency are outlined below. First, all calibrated values that emerge from the analysis program must be divided by 2 because of the structure of a double-switched flux measurement, where the source position is measured in the ANT channel first, but is then measured in the REF channel during the ‘OFF’ integrations. This means that the results of a flux procedure contains twice the actual source value.

Further, a small correction must be applied to correct for the total power difference during observations and during calibration procedures (when the system points to the internal loads for stability). As a result, the non-linearity probed in the two cases is slightly different, and is expressed as $\kappa_{NL} = \eta(P_{CAL})/\eta(P_{FLUX})$. In principle, this correction depends on zenith angle because of varying power levels, but this dependence is well below one percent. The effective correction is 1.015 for the 1992 season and 1.0025 for the 1993 season. The resulting values are now real antenna temperatures.

Of course, antenna temperature differences differ from sky temperature differences by the beam efficiency. In section 3.3.4, this was determined to be $\varepsilon_b = 0.740 \pm 0.039$, the error being a systematic scaling error without effect on the statistical significance of the result.

The last multiplicative correction is caused by the noticeable departure of the microwave background spectrum from the Rayleigh-Jeans approximation at our observing frequency.

The correction factor is $\kappa_\nu = \Delta T_0 / \Delta T_\nu$ (section 2.1) and converts the measured temperature difference to the one at zero frequency. At our center frequency of 32GHz, $\kappa_\nu = 1.054$.

The corrected SZ temperature difference therefore becomes

$$\Delta T_{corr} = \Delta T_{raw} \frac{e^{\tau \chi_{atm}}}{2\mathcal{E}_B} \kappa_{NL} \kappa_\nu.$$

5.3.2. Results

The results have been calculated as weighted means. While the thermal editing procedures result in a restricted range of data pair error bars, it is still apparent that remaining data with larger error bars reflect generally worse observing conditions. For this reason, an unweighted mean was unacceptable. Fortunately, the choice of averaging method makes little difference to the result; this is explored in the section on data quality below.

In addition, some weighting must be employed to account for the variable integration times at different zenith angles. Standard error weighting takes this into account very naturally, further justifying the choice of that method.

Table 5.3 shows the overall results by position in the cluster and by observing season. The values are zero-frequency SZ temperature differences, and represent the effective observed SZ effect for each position. All the corrections mentioned previously have been applied—these values can be compared directly to the predictions of the X-ray model, once corrections for discrete sources have been applied. N denotes the number of data pairs used in the calculation.

TABLE 5.3
OBSERVED SZ DECREMENTS

Position	Season	N	ΔT_0
C	1992	1518	$-261.9 \pm 33.2 \mu\text{K}$
	1993	955	-269.8 ± 60.8
	overall	2473	-263.7 ± 29.2
N	1992	583	-47.0 ± 64.0
	1993	295	-128.9 ± 81.7
	overall	878	-78.2 ± 50.4
CS	1993	291	-17.6 ± 71.2

NOTE values are zero-frequency SZ temperature differences. The effect of discrete radio sources has not been removed.

From the table it is clear that by far the most and best data have been gathered for the center position in the cluster, C. In terms of raw antenna temperature, the effective 1σ uncertainty

in that position corresponds to less than $20\mu\text{K}$! The overall value in position C has a significance of more than 9σ .

This error does not yet take into account the systematic scaling error in the beam efficiency. Since it is constant for all observations, it should not be used to judge the statistical reliability of the result. In addition, this error is small: even for the overall data on position C, the noise is increased only from $29.2\mu\text{K}$ to $32.4\mu\text{K}$.

It is also encouraging that the two seasons agree extremely well in the position with the best data; even the slight discrepancy between the two seasons in position N is statistically insignificant: the difference of $81\pm 104\mu\text{K}$ is not even an 0.8σ effect.

Even though the final interpretation of these values depends on the X-ray model results discussed below, it is nonetheless encouraging to see the values in positions N and CS as close to zero. This should enhance our confidence that this experiment successfully controls systematic effects, in particular differential ground pickup.

Finally, it is remarkable to see the observed values for the effective SZ effect in Coma as small as these. It will be shown below that these values are entirely within the expected range, but it becomes understandable why the effect was not detected ten years ago!

5.4. TESTS OF DATA QUALITY

Below are the results of several tests of the quality of the data. Overall, the data pass several important tests very well, but some tests (such as stability with parallactic angle) show inconsistencies with marginal statistical significance.

5.4.1. *Averaging Procedures*

For the results above we have used weighted means, with some justification. Fortunately, the results change only very little if we use unweighted means, as is shown in table 5.4. This robustness is very encouraging.

In principle, it would also be possible to calculate a mean based on the maximum of a Kolmogorov-Smirnov (KS) probability in the space of mean and standard error.

TABLE 5.4
COMPARISON OF WEIGHTED AND UNWEIGHTED MEANS

Position	Season	ΔT_0		Difference
		weighted	unweighted	
C	1992	$-261.9 \pm 33.2 \mu\text{K}$	$-247.3 \pm 35.7 \mu\text{K}$	$-14.6 \mu\text{K}$
	1993	-269.8 ± 60.8	-267.4 ± 60.8	-2.4
	overall	-263.7 ± 29.2	-252.4 ± 30.8	-11.3
N	1992	-47.0 ± 64.0	-38.1 ± 66.5	-8.9
	1993	-128.9 ± 81.7	-138.5 ± 83.3	9.6
	overall	-78.2 ± 50.4	-77.2 ± 51.9	-1.0
CS	1993	-17.6 ± 71.2	-2.4 ± 72.0	-15.2

NOTE see remarks for table 5.3.

However, while the maximum is extremely well-defined in the standard error direction (typically yielding errors slightly less than those from conventional calculations), the maximum is usually very broad and noisy in the mean's direction. While this is a reflection of the internal data uncertainty, this makes it impractical to derive a good estimate of the mean from this method.

5.4.2. Subsets

A standard way to assess data quality is to subdivide the full data set into several subsets and to compare the consistency of the result among them. One obvious way to subdivide the data is by leading and trailing reference field; this has been done in table 5.5. Note that the values given there are still fully ground-subtracted; the subdivision merely distinguishes between the two directions in which the ground subtraction takes place.

TABLE 5.5
LEADING AND TRAILING SUBSETS

Position	Season	ΔT_0	
		Main-Leading	Main-Trailing
C	1992	$-240.8 \pm 53.5 \mu\text{K}$	$-273.2 \pm 42.2 \mu\text{K}$
	1993	-226.6 ± 89.7	-298.6 ± 81.7
N	1992	-80.3 ± 90.8	-13.8 ± 89.2
	1993	-161.7 ± 154.5	-113.7 ± 95.3
CS	1993	-170.5 ± 97.7	$+144.9 \pm 103.3$

NOTE see remarks for table 5.5.

Clearly, the subsets for position C are entirely consistent with each other, and so appear the subsets for position N, once the larger error bars are taken into account. The two subsets for position CS are not very consistent at all. It is not clear how significant that difference is, given the small amount of data, taken during generally unfavorable weather conditions. However, the consistency of the 1993 data on position N is good—the data for position CS were taken during the same time period. It could well be that this slight discrepancy in the two subsets for position CS is caused by discrete sources in one of the reference fields. This should become clear once the entire discrete source data set is available.

A further possibility of subdividing the data set is by the position of the OFF beams on the sky—this binning by parallactic angle is done below.

5.4.3. Ground Subtraction

The quality of the ground subtraction technique is difficult to assess, since the expected contamination is small. However, the success of our subtraction of differential ground spillover is illustrated well in figure 5.2.

While the error bars in the bottom panel of figure 5.2 are larger than in the top one (this is expected because of the subtraction), the improvement because of ground subtraction is clearly visible. The data in the bottom panel show no statistically significant deviation from a uniform mean, while those in the top panel clearly exhibit non-random trends. This should be taken as confirmation that residual differential ground spillover has been subtracted properly with our observing technique.

5.4.4. Statistical Distribution

One of the main aims of all the various levels of controls, subtractions, and editing methods is to produce a data set whose value distribution resembles a Gaussian distribution as closely as possible. Apart from looking at the distribution by eye, one can use the Kolmogorov-Smirnov (KS) test to judge the departure from any model distribution. Given the correct model distribution, the test will then produce the probability $P_{KS}(D > D_{obs})$ of a chance occurrence of the maximum cumulative departure in the data, D_{obs} , or any larger

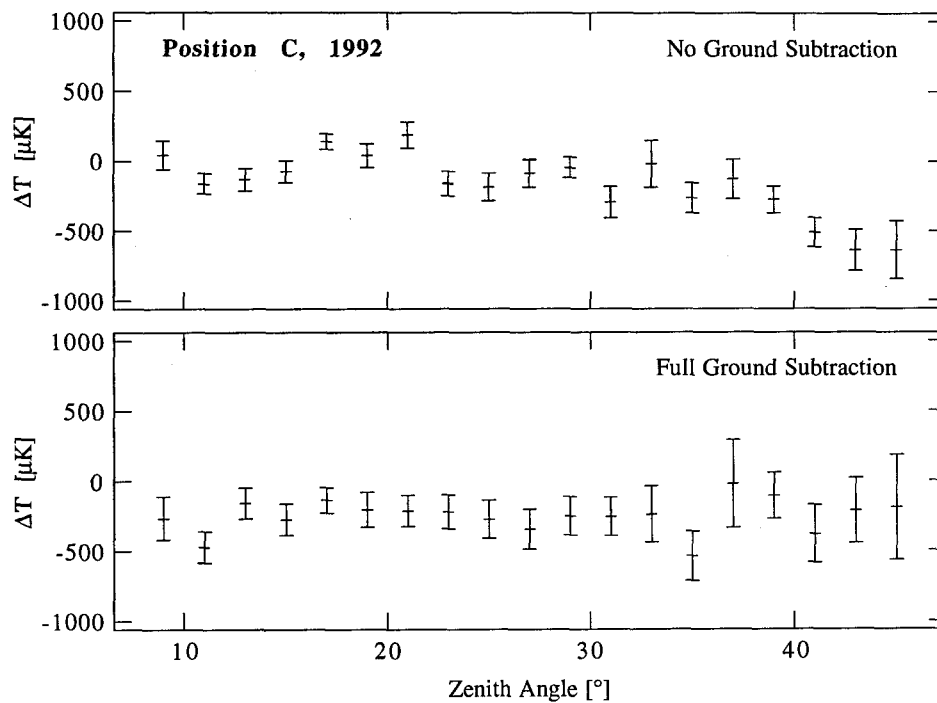


Figure 5.2 The effect of our technique to subtract differential ground spillover. The top panel shows flux data binned in zenith angle on the Main field only, without Leading or Trailing reference fields. The bottom panel shows the same data, with the reference fields subtracted. The data set used here is the statistically fullest one: 1992 observations of position C in Coma. The temperature units are zero-frequency SZ temperature differences.

value. In other words, this probability can be thought of as the probability of agreement between the actual distribution with the model one.

When this KS test is used at first, probability values of less than 1% are often produced. This illuminates the trick of the whole matter: which model distribution do we use to judge the data? In principle, we want a Gaussian one, but which values do we use for the parameters of mean and standard deviation? With data sets like the ones in this experiment, the value of the probability P_{KS} is not very sensitive to the mean, but it depends very strongly on the standard deviation, since most of the data's departure from a Gaussian occurs in the wings of the distribution. In my experience, a 1% error in the model's standard deviation can reduce the probability by 10% or more.

For this reason, fitting algorithms are commonly used that vary the input model parameters until a maximum value of the probability has been reached. This is an entirely legitimate

procedure, since it represents merely a slightly different weighting scheme from the ones commonly used. However, this method has not been used in the calculation of any of the results, because of its insensitivity to the mean value of the data, which is of primary interest to us. The presence of local minima can often mislead the algorithm severely, since the minimum is so shallow and broad.

However, this method can indeed be used to judge the quality of the data set and to find the maximum KS probability, since it utilizes an entirely legitimate model of the data. As long as it is kept in mind that this maximum KS probability value does not correspond to the mean and standard deviation derived by other means, its use should not be misleading.

The KS probability values were calculated for all data sets and many of their subsets. In general, the probabilities were between 40% and 80% for the full data sets that are divided only by season and position. The larger the data set, the smaller the KS probability value, perhaps reflecting the fact that structures within the data are more statistically prominent in large than in small data sets. This conclusion of inherent structure in the data is also supported by the fact that sub-groups of these large sets, *e.g.*, divided by parallactic angle, all have much larger probability values, between 70% and 95%.

5.4.5. Variability with Parallactic Angle

We can form further subsets of the data by dividing them into ranges of the parallactic angle ψ . This angle corresponds to a linear track on the sky, where bins of 18° size are roughly independent because of the beam's size. In principle, there should be no variability in the data except for the presence of discrete sources and the presence of actual and legitimate structure in the observable SZ effect in the cluster itself. Since both of those effects are most likely on the order of only a few μK in the case of Coma, given our observing arrangement (discussed below), variations in the data with parallactic angle could point to problems with the data.

In fact, the data are not entirely constant with parallactic angle, but at a statistically marginal level. Figure 5.3 shows the parallactic angle distribution of our statistically most descriptive data set, the 1992 data on position C. The figure also illustrates the strongly non-uniform distribution of the data with parallactic angle, a result of Coma's declination, which causes it to transit very close to the zenith at OVRO.

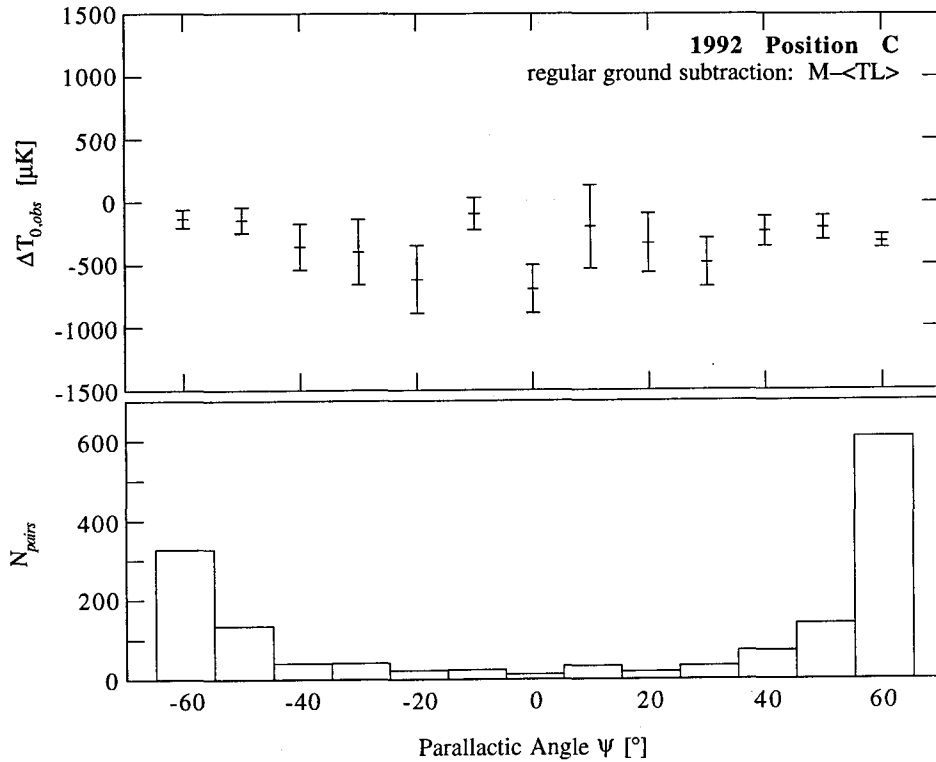


Figure 5.3 The top panel contains the fully ground-subtracted 1992 data set on position C, binned by parallactic angle in 10° bins. The units are zero-frequency SZ temperature differences. The bottom panel shows the number of complete data pairs in each bin; this explains the large error bars in the middle of the top pane.

The inner parallactic angle bins in figure 5.3 should be disregarded; some of them only possess 20 or fewer complete data pairs, as can be seen in the figure's bottom panel. However, the difference between the lowest and the highest parallactic angle bin is noticeable. The numbers from those two bins are listed for all the data sets in table 5.6.

TABLE 5.6
PARALLACTIC ANGLE DEPENDENCE

Position	Season	ΔT_0		Difference	ΔT_0
		<i>B</i> : $\psi < -55^\circ$	<i>T</i> : $\psi \geq +55^\circ$	<i>T-B</i>	<i>M</i> : $-55^\circ \leq \psi < +55^\circ$
C	1992	$-130.5 \pm 71.3 \mu\text{K}$	$-317.0 \pm 54.3 \mu\text{K}$	$-186 \pm 90 \mu\text{K}$	$-274.0 \pm 51.9 \mu\text{K}$
	1993	-274.6 ± 112.1	-402.7 ± 78.5	-128 ± 137	-101.7 ± 129.7
	overall	-172.1 ± 65.3	-344.7 ± 44.7	-173 ± 79	-250.2 ± 59.4
N	1992	$+123.2 \pm 151.6$	-103.8 ± 88.4	-227 ± 176	-80.3 ± 111.9
	1993	-20.8 ± 112.9
	overall	$+30.6 \pm 90.5$	-103.8 ± 88.4	-134 ± 126	-80.3 ± 111.9

NOTE see remarks for table 5.3.

Clearly, the trend visible from figure 5.3 is confirmed by the binned data. The ‘bottom’ parallactic angle bin has a consistently higher temperature than the ‘top’ bin, while the ‘middle’ bin falls in between the two. The fact that the difference between the top and bottom bins is not zero, but in fact always negative and similar in size regardless of observing season and observing position is worrisome, and no real explanation has been found yet for this behavior.

On the other hand, while this discrepancy is certainly suggestive, it is not statistically significant. The overall $T-B$ difference for the C position is less than a 2.2σ value, and even the combination of the overall differences in both observing positions, C and N, $-162 \pm 67 \mu\text{K}$, barely tops the 2.4σ mark.

Only with much improved data will it become possible to ascertain a bias and to find its causes, but this aspect of the data needs to be monitored in the future.

It is possible, however, to establish that this parallactic angle variation is not an artifact introduced by the reference fields. Figure 5.4 shows the same data as in figure 5.3, but rather the difference between the leading and the trailing data subsets binned by parallactic angle. No structure is apparent in that difference, and the well-defined end points of the parallactic angle range are both consistent with zero.

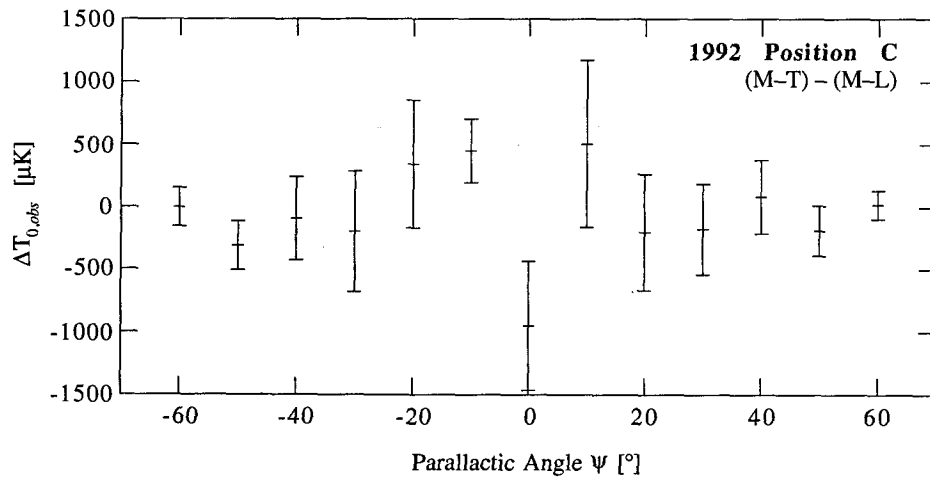


Figure 5.4 Difference between the trailing and leading reference field subsets in the 1992 data set on position C, binned by parallactic angle. Each term entering the subtraction is fully ground subtracted by itself—this graph shows the reliability of this subtraction as a function of parallactic angle.

6. DISCRETE RADIO SOURCES

Discrete radio sources are the dominant astronomical effect that can corrupt the measurement of the SZ decrement (see *e.g.*, Schallwisch and Wielebinski 1979). Several strategies are usually chosen to reduce their impact, and all have been employed in this experiment. First, one chooses to observe a cluster in which the flux density of discrete radio sources is small. For this reason, for instance, the cluster A426 was not observed, since it contains the well-known radio source 3C84. The Coma cluster is relatively ‘clean’, with no dominating radio sources.

Second, one chooses an observing frequency at which the flux density of most radio sources is small. For the standard synchrotron source, the majority of objects at ‘low’ frequencies (below 60GHz), the flux density falls rapidly with frequency in a power-law fashion $S_\nu \propto \nu^\alpha$, with α typically around -0.75 . For this reason, the higher the observing frequency, the smaller the influence of discrete radio sources with steep synchrotron spectra. At mm and sub-mm frequencies, the influence of some black-body sources becomes noticeable and must be taken into account. Since we are far from that region, it is disregarded here.

Third, a large beam size is advantageous because of the lower point-source gain of the telescope. This affects not so much the average total flux in the telescope’s beam, since the number of radio sources in the beam increases at the same rate that the gain decreases; the total flux at various beam sizes depends on the shape of the source count function. The advantage works differently. The bulk of radio source contamination usually comes from a few bright sources in the telescope’s beam. With lower point source gain, those sources that affect the measurement most strongly are much brighter than with a small beam. As a result, efforts to measure those source fluxes will be much simpler to accomplish. In addition, the mean background source contamination is removed by our differencing technique. This mean source contamination (mostly from faint sources) is much better measured by a large beam.

After taking steps to minimize the influence of discrete sources, corrections for them become feasible. If one knows where to find the sources, one could in principle measure their flux densities with a high-gain telescope (such as the VLA or the OVRO 40m tele-

scope). That is precisely the difficulty, however. Because of many difficulties, no sensitive radio source surveys exist at frequencies as high as 32GHz.

Instead, one refers to lower-frequency surveys and extrapolates them to higher frequencies. Unfortunately, the completeness of the radio source subtraction at the SZ frequency of 32GHz is not at all guaranteed. Since we have almost no information about the spectral distribution of faint sources at high frequencies, selection of sources at lower frequencies and assuming the above spectral index can lead to problems. Ideally, to find all sources that could possibly affect our high-frequency measurements, we would select discrete sources at a low frequency (so that they can be found efficiently) with high enough sensitivity so that sources with any reasonable spectral index will be found.

With this in mind, the best way to accomplish the source subtraction is with a sensitive survey of the whole region using the VLA at low frequency. It would include the Coma region as well as the Leading and Trailing reference regions. I am currently working on this project, and its results should become available within the next year. Objects selected from this low-frequency survey could then be monitored with the OVRO 40m telescope at 18.5GHz and the OVRO mm-array at 32GHz to measure their high-frequency flux densities and their variability.

Before these results are available, however, sources were selected with lower sensitivity from existing surveys. These sources were then observed with the OVRO 40m telescope at 18.5GHz, very close to the actual frequency of the observations. Extrapolation to the SZ frequency of 32GHz is made on this basis.

This selection has been made with the central SZ position in mind, since it is used for the calculation of the Hubble constant below. However, no comparable data such as the 5C4 survey were available for the leading and trailing reference fields.

6.1. EXISTING SURVEYS

Two of the most sensitive existing surveys are the 5C4 survey of the Coma region (Willson 1970) and the Condon 5GHz survey (Condon *et al.* 1989).

The 5C4 survey was made at 408MHz with the Cambridge one-mile telescope in 1968, achieving a flux density limit of 20mJy for the area in question. The Condon survey was

made at 4.85 GHz with the late Green Bank 300-foot telescope in 1987. The object catalog was produced from the survey data by Becker *et al.* (1991), and has a detection flux limit of 25mJy for the area in question. Because most objects in the region at these flux levels are steep-spectrum synchrotron sources, the 5C4 survey provides the far more stringent flux density limit. On the other hand, the similarity of the two flux limits allows the detection of flat-spectrum sources of sufficient brightness.

Given the flux limit of the 5C4 survey and our experiment's requirement of $T_{A,src} \leq 10 \mu K$, all sources with a spectral index of $\alpha < -0.51$ will be found in the survey. This is adequate for finding standard steep-spectrum sources. Only sources with a flatter spectrum will escape detection by this method. While this is certainly not a permanent solution, it qualifies as a respectable first cut.

From the 5C4 survey, a group of 21 sources was identified for further investigation at higher frequency; only 12 sources in the Condon survey can be found in the general area. From these, only 7 were of immediate concern to the data, and all have counterparts in the 5C4 survey. The 18.5 GHz measurements of this group of 21 candidates are described below.

6.2. OBSERVATIONS AT 18.5 GHZ

These observations were carried out from April 23, 1992 to May 1, 1992 using the OVRO 40m telescope with its microwave background K-band receiver. This receiver is very similar to that of the 5.5m telescope, except that it lacks some monitoring and calibration capabilities. Its first stage is a cryogenically cooled HEMT amplifier with a receiver temperature of about 40K. In contrast to the 5.5m telescope's receiver, this is a heterodyne receiver.

The complement of pre-selected sources was observed completely for several nights, so that acceptable statistics could be found. Using the 40m telescope's standard calibration along with observations of a calibration source (3C286), the flux densities shown in table 6.1 were obtained. This table also contains the lower-frequency fluxes, the spectral indices and the fluxes at 32GHz. In addition, for each source the fraction of its flux seen by the 5.5m telescope's beam is listed.

The effect of sources in the reference arcs is minor indeed; given the relative integration time spent in the applicable parallactic angle range and the fact that each reference beam has

TABLE 6.1
DISCRETE SOURCE FLUX DENSITIES

Name	Position (1950)		$S_{\nu,obs}$		α	$S_{\nu,extr}$	$T_{A,eff}$		
	α	δ	1.4GHz	18.5GHz			32GHz	center	ref
5C4.53	12 55 19.9	+28 13 49	22 ± 4	< 2
5C4.54	12 55 27.5	+28 22 18	60 ± 17	16 ± 2	-0.51 ± 0.12	12.1 ± 5.7	...	4.0 ± 1.1	+23
5C4.70	12 56 17.2	+28 10 19	23 ± 3	< 2
5C4.73	12 56 29.3	+28 34 02	64 ± 3	6 ± 1.5	-0.92 ± 0.10	3.6 ± 1.1	...	3.0 ± 0.9	+60
5C4.74A	12 56 32.7	+28 03 00	27 ± 8	< 2
5C4.74B	12 56 36.0	+28 02 50	28 ± 8	< 2
5C4.77	12 56 36.1	+28 27 50	10 ± 2	< 3
5C4.81	12 56 56.7	+28 10 50	367 ± 29	24.5 ± 2	-1.05 ± 0.04	13.7 ± 2.2	4.4 ± 0.7
5C4.85	12 57 10.7	+28 13 48	178 ± 19	24 ± 1.5	-0.78 ± 0.05	15.5 ± 2.9	19.7 ± 3.7
5C4.92	12 57 42.0	+27 45 48	54 ± 5	10 ± 1.5	-0.65 ± 0.07	7.1 ± 1.6	...	9.9 ± 2.2	+84
5C4.95	12 57 50.0	+27 55 41	19 ± 2	6 ± 1	-0.45 ± 0.08	4.6 ± 1.2	...	1.6 ± 0.4	+74
5C4.102	12 58 02.0	+28 25 33	10 ± 2	< 2
5C4.105	12 58 04.0	+28 46 21	155 ± 9	210 ± 15	+0.12 ± 0.04	226 ± 28	...	< 0.1	-79
5C4.114	12 58 26.3	+28 24 13	49 ± 2	< 4
5C4.122A	12 58 38.7	+28 34 14	40 ± 15	< 2
5C4.122B	12 58 46.5	+28 34 39	62 ± 17	4 ± 1	-1.06 ± 0.14	2.2 ± 1.2	...	0.4 ± 0.2	-55
5C4.125	12 58 48.6	+28 30 48	16 ± 3	< 2
5C4.127	12 58 45.6	+28 37 49	118 ± 6	4.5 ± 1	-1.27 ± 0.09	2.2 ± 0.6
5C4.128	12 58 56.2	+27 53 33	62 ± 5	5.5 ± 1	-0.94 ± 0.08	3.3 ± 0.8	...	3.9 ± 0.9	+43
5C4.129	12 58 57.0	+28 23 37	37 ± 3	8 ± 1	-0.59 ± 0.06	5.8 ± 1.2	...	12.8 ± 2.7	-33
5C4.140	12 59 36.4	+28 18 56	31 ± 4	< 2

Notes. the 1.4GHz fluxes come from the 5C4 survey (Willson 1970), while the 18.5GHz fluxes come from observations with the OVRO 40m telescope. The spectral index α is produced by these two fluxes and is used to produce the extrapolated flux densities at 32GHz. The last three columns illustrate the actual antenna temperatures of these sources when observing the C position in Coma with the 5.5m telescope. The *center* values are the effect on the central beam, while the *ref* values are the effect on one of the reference beams if the telescope observes at parallactic angle ψ .

All flux densities are in mJy, and all antenna temperatures in μK , with those smaller than $0.1\mu\text{K}$ denoted by ...

only half weight, their influence is negligible. This is particularly true since the parallactic angle range of the observations cannot exceed the range $-65^\circ < \psi < +65^\circ$. This excludes the large contribution from 5C4.92. For these reasons, the discrete source contributions from the reference beams have been disregarded completely, with a maximum error in antenna temperature of less than $0.5\mu\text{K}$.

Mercifully, the only flat-spectrum source in this collection, 5C4.105, is far enough removed from any of the beams that it does not affect the experiment.

On the other hand, the antenna temperature contributed by 5C4.81 and 5C4.85 to the center beam patch must be taken into account. Its effect is to reduce the size of the apparently visible SZ decrement.

No variability information is available for these sources. However, all sources of concern have steep radio spectra, which make variability unlikely. In addition, 5C4.81 and 5C4.85 are extended steep-spectrum objects, which are not variable. Therefore, the error in the known source correction from variability is disregarded.

6.3. THE HALO SOURCE

The very center of the Coma cluster harbors a monstrous central halo source, an extended area emitting very steep-spectrum synchrotron radiation. The origin of such halo sources is disputed, but their age is apparent from the steepness of their synchrotron spectra. There are about 5 known cluster halo sources, all in hot, regular clusters; the most recent one was detected in the hottest known cluster, Abell 2163 (Herbig *et al.* 1993).

The contribution of this halo source to the antenna temperature of the 5.5m telescope in the C position is very small because of the halo's steep spectral index, but it can be included easily. From the low-frequency image of the cluster (Venturi *et al.* 1990), we find that the beam size of the 5.5m telescope is about 22 times larger than the WSRT beam used for the low-frequency observations at 326MHz. Since the average surface brightness of the halo source is about 10mJy/beam in the image, the total surface brightness seen by a beam of the size of the 5.5m telescope at 326MHz is about 220mJy.

Given the halo source's spectral index of $\alpha = -1.3$ (Venturi *et al.* 1990), the total surface brightness in the 5.5m telescope's beam at 32GHz is 0.56mJy, or 2.6 μ K. Given the approximate nature of this analysis, the error estimate may be as large as $\pm 1\mu$ K. This halo source affects the central beam only; the reference beams do not fall within its extent.

6.4. CORRECTED RESULTS AND UNDETECTED SOURCE ALLOWANCE

Since we apply several factors to the antenna temperature of the SZ observations, these same factors must be applied here before the value of the correction can be applied to the total zero-frequency SZ effect from section 5.3.2 above. Excluding the factor of 2 which does not apply to these measurements, the effective antenna temperatures of the discrete sources of table 6.1 must be multiplied by 1.439.

For the discrete sources, the total correction therefore becomes $34.7 \pm 5.5 \mu\text{K}$ in the units of the total zero-frequency SZ effect. Correction for the halo source increases this by another $3.7 \pm 1.6 \mu\text{K}$ to a total correction of $38.4 \pm 5.7 \mu\text{K}$.

Unfortunately, this is a correction only for the effect of those discrete sources that we have found in the 5.5m telescope's beam with our preliminary selection. Most likely, there are further sources that have not been found, but that collectively affect the experiment's result. To correct for this effect is very difficult, but it is plausible that the sum of these additional sources will not exceed the discrete source correction actually made. Therefore, an additional value of $0 \pm 35 \mu\text{K}$ is added to the result. It is arguable whether this error bar should be symmetric (it should not if most sources are associated with the cluster), but there is no real substitute for awaiting the much more definitive source subtraction based on our VLA data.

The final corrected value for the observed SZ effect in position C in the Coma cluster therefore becomes $-302.1 \pm 32.9 \mu\text{K}$ without an undetected discrete source allowance, and $-302.1 \pm 48.0 \mu\text{K}$ with the conservative assumption of such an allowance. These errors contain the uncertainty of the beam efficiency. It is clear that the final VLA discrete source data set will help to reduce the error estimate in this number substantially.

7. INTERPRETATION

Previous chapters have established the observed value for the SZ effect in the Coma cluster, observed with a particular pattern on the sky, centered on three distinct positions, and corrected for all known effects that could affect that value. So far, this accomplishes only one unqualified astrophysical conclusion: that the cosmic microwave background must indeed issue from behind the Coma cluster of galaxies. While that is a noteworthy result, it is not new and is much better argued by SZ observations in distant clusters of galaxies (Birkinshaw *et al.* 1993).

Most useful results from these observations rely on their combination with X-ray observations of the same cluster gas. Some of the stated goals of SZ observations cannot be met with the data obtained for this thesis—inferences for the intracluster medium require much more complete SZ profile data than we possess. In the future, this aspect will become far more tractable as profiles and maps of the effect will become more available.

On the other hand, the data that are available do allow a measurement of the Hubble constant. Far from permitting a self-consistent profile solution, this measurement relies on the quality of existing X-ray gas models, but this is precisely the advantage of Coma over the clusters whose SZ effects have been observed to date. Below I will give a short portrait of the Coma cluster in X-rays and trace the development of its X-ray models. This will be followed by an overview of the distance derivation, and the application of X-ray models to the data for a measurement of the Hubble constant.

7.1. THE COMA CLUSTER IN X-RAYS

7.1.1. *Description*

Recently, a wealth of new and more accurate information on Coma's X-ray atmosphere has become available. Both all-sky survey and pointed images from the Rosat X-ray telescope are much improved over the best images taken with the Einstein X-ray telescope.

Even with such improved data, the structure of the Coma cluster in X-rays still has one striking feature: it is remarkably regular, and, apart from a certain amount of apparent ellipticity, is very symmetrical. On the basis of newly available Rosat data (*e.g.*, White *et al.* 1993), it becomes arguable that the origin of this ellipticity is due to the distorting effect of sub-clumps in the cluster's atmosphere around the cluster's dominant galaxies. Still, the overall departure of the cluster from regularity is minor; arguments emphasizing such departures should be seen in the context of our traditional view of the prototypical nature of the Coma cluster.

The cluster's atmosphere also possesses a long gas extension to the southwest—this culminates in a broad surface brightness peak about 1° away from the center of Coma, roughly coinciding with the position of the extended radio source 1253+275 and the galaxy NGC 4839. It is quite plausible to suspect that this extension represents a subcluster that is seen in the process of merging with the main part of the cluster. The gas density and the temperature of this component is probably small compared to Coma itself; an Exosat pointing near the position of this extension gives a temperature of about 1.9keV (Edge 1989). For this reason, the extension has been excluded completely from the model calculations below.

The increasing sophistication and realism of our view of the Coma cluster's gas atmosphere, among other clusters, is one of the most important advances for the interpretation of the SZ effect, even though it does not simplify this interpretation. Currently, such an enhancement is just at the horizon, but more work is needed to establish a consensus on how best to translate a more realistic view of the cluster atmosphere into an acceptable three-dimensional model.

In part, such a translation requires a better understanding of the cluster's temperature structure, which relies on data that have not yet undergone the same quality improvement as the soft X-ray imaging observations. Proposed pointed observations with the ASCA satellite offer one of the best chances to accomplish this in the coming few years.

In the meantime, however, we must resort to the simplified approach of a spherically symmetrical cluster atmosphere whose two main physical parameters, electron temperature and density, can be described by relatively simple radial models.

7.1.2. X-ray Models

The first models of Coma's X-ray atmosphere followed basic assumptions about the cluster's temperature structure and were generally expressed as standard beta models with

$$n_e(r) = n_{e,0} \left[1 + \left(\frac{r}{r_{core}} \right)^2 \right]^{-\frac{3}{2}\beta} .$$

This choice of a density structure is physically reasonable, since it describes the expected structure from an equilibrium sphere of gas, in which the (isothermal) velocity dispersion of the cluster's member galaxies is expressed in terms of the gas temperature in the dimensionless beta parameter as $\beta = \mu m_p \sigma^2 / kT_e$, given a polytropic dependence of the electron temperature (Cavaliere and Fusco-Femiano 1976, 1978; Bahcall and Sarazin 1978).

As to the temperature distribution of the gas, these early models either assumed an isothermal structure (with an implied, but unspecified cutoff at some large radius) or a polytropic structure with

$$T_e(r) = T_{e,0} \left[\frac{n_e(r)}{n_{e,0}} \right]^{\gamma-1} .$$

Of course, a polytrope with $\gamma = 1$ is an isothermal sphere, while a polytrope with $\gamma = 5/3$ is an adiabatic distribution.

Initially, such models fit the data quite well, but it became apparent with a careful analysis of Einstein's imaging data combined with accurate spectral information from the Tenma X-ray telescope, that while an isothermal core was required for the cluster, the fit to a polytropic model at large radii was best (Hughes *et al.* 1988a). In fact, both a pure isothermal and a pure polytropic model could be ruled out to high confidence from those results. Instead, a hybrid model was established that forms the basis of the model used in the interpretation of this experiment.

This hybrid model combines a central, isothermal core with a polytropic decrease of the electron temperature at large radii, while retaining a beta-model dependency for the electron density (Hughes *et al.* 1988b):

$$T(r) = \begin{cases} T_{iso} & r \leq r_{iso} \\ T_{iso} \left[\frac{1 + (r/r_{core})^2}{1 + (r_{iso}/r_{core})^2} \right]^{-\frac{3}{2}\beta(\gamma-1)} & r > r_{iso} \end{cases}$$

In addition to the traditional parameters of central electron density $n_{e,0}$, core radius r_{core} , and β , this hybrid model also uses the central isothermal temperature T_{iso} , the radius to which the isothermal model holds r_{iso} , and the polytropic index γ . From binding mass considerations, γ is restricted to the range from 1 (the isothermal case) to about 1.5.

Such an approach was taken initially with Einstein data on the cluster, concentrating on a main 'North' field as well as the average of all IPC fields (Hughes *et al.* 1988b). The radius to which these fits could be made was about 20', except for the 'North' field, which reached to a radius of 40'.

This restriction in radius is problematic for a cluster of Coma's size, which has considerable surface brightness at such large radii. The field of view of most X-ray telescopes is not sufficient to fit Coma comfortably, and radius restrictions of the models were therefore unavoidable. This situation has changed with the all-sky data from Rosat, which provides scanned data to arbitrarily large radii. In their analysis of the Rosat all-sky data on Coma, Briel *et al.* (1992) measure the cluster's surface brightness to a radius of 100'. However, the low energy range of Rosat does not allow a temperature measurement, so that this parameter cannot yet be determined from the large-radius observations. Briel *et al.* fit a beta-model of electron density to the surface brightness data, and achieve a very good fit. Given the results of the Tenma data (Hughes 1988a), an isothermal temperature distribution is very unlikely to hold at large radii. Even though the Rosat all-sky surface brightness was fit with an isothermal model, a hybrid temperature model similar to the one above was actually used in the modeling to satisfy the much more stringent temperature constraints of the Tenma data.

Even further data are available in the form of long-integration pointed observations of several fields in and around the Coma cluster with Rosat (White *et al.* 1993), but while these data greatly enhance our appreciation of the details of the cluster's structure and its sub-units, no radial density model has yet been produced from those data as a compromise

between the recognition of the cluster's morphological structure and the expediency of an overall model of the cluster atmosphere's physical parameters. Unfortunately, this means that none of the high-quality pointed data on Coma are used in this thesis; yet the uncertainty in the X-ray atmosphere's model is probably the most important contribution to the systematic uncertainty in our derivation of the Hubble constant.

Clearly, these models will have to be revisited with the inclusion of the pointed Rosat surface brightness data as well as with forthcoming temperature data from ASCA.

In all of these models, the apparent ellipticity of the cluster's atmosphere has been disregarded and the data were azimuthally averaged. Given the tenuous nature of some of the model fits, such a simplification is an entirely reasonable compromise. Most likely, it will be overcome by the deeper structural analysis that is likely to result from recent data. It will be shown below that an initial adaptation of the hybrid model to an elliptical cluster shape affects our derived value of the Hubble constant by no more than 10%.

7.2. DERIVATION OF THE HUBBLE CONSTANT

7.2.1. Theory

An excellent synthesis and generalization of the various methods to derive the Hubble constant from the combination of SZ and X-ray data is given in Birkinshaw *et al.* (1991). Since the data in hand can be treated in a much more simple-minded fashion, this method is summarized in appendix F. It uses an angular coordinate system with angular radius θ from the cluster center, azimuthal angle ϕ about the line of sight, and length along the line of sight ζ , expressed in angular units by $\zeta = l/D_A$, dividing by the angular diameter distance.

The interpretation of this experiment uses a shortcut implicit in the derivation of the Hubble constant. If one derives an X-ray model that predicts radial runs of electron temperature and density, the observed X-ray surface brightness distribution is in fact contained in the model. In addition, the normalization of the X-ray model was undertaken with a particular assumption of a value of the Hubble constant.

The inversion is straightforward: given an X-ray gas model, it can be used to predict the SZ decrement for a particular line of sight at some angular distance θ from the cluster center (assuming azimuthal symmetry). In principle, this value is then compared to the observed SZ decrement. If the two values do not agree, the method concludes that the error is due to an incorrect ‘trial’ value of the Hubble constant. The implied value for the Hubble constant then becomes

$$\frac{H_{0,obs}}{H_{0,model}} = \left(\frac{\Delta T_{model}}{\Delta T_{obs}} \right)^2,$$

with the ΔT 's being the usual SZ temperature change at zero frequency. This relation also points out the unfortunate dependency of the constant's measurement on the square of the observed and the predicted SZ temperature decrements.

In the X-ray models quoted, a value of $H_{0,model} = 50 \text{ km s}^{-1} \text{ Mpc}^{-1}$ was used throughout.

7.2.2. X-ray Model Predictions

Basic Method

The calculation of the SZ temperature prediction from the X-ray gas model is straightforward, but does involve some complications. First, a Runge-Kutta integration with an automatically variable step size (see Press *et al.* 1986) is carried out along a line of sight offset from the cluster center by an angular radius θ . This integration is carried out to a distance of 200' along the line of sight ζ from the sky plane intersecting the cluster center (for a total line-of-sight path of 400'). The accuracy of the integration is about 10^{-5} . This gives us a pencil-beam SZ decrement as a function of distance from the cluster center.

To compare the prediction to the actually observed value, another transformation must be accomplished: given the shape of the gas atmosphere, what would we expect to see with a telescope with our beam size and our beam separation. To achieve the first, we add the product of the beam weighting (the sum of the beam is normalized to 1, since the observed value already includes a correction for the forward beam efficiency) and the underlying SZ effect over a 41×41 step grid with a 0.45 step size. The resulting patch is calculated 3 times: once for the central ON beam and once for each of the two subsequently averaged

and subtracted OFF beams. The result from this subtraction represents the effectively expected observed SZ temperature change at zero frequency using our observing configuration, $\Delta T_{0,eff}(\theta)$.

This predicted value can be calculated as a function of the angular distance from the cluster center, but also over a range of parallactic angles at each individual value of θ .

Specific Models

Several models were used to calculate the predicted SZ decrement. First, there are the two hybrid models of Hughes *et al.* (1988b), covering the ‘North’ field as well as the average of all IPC fields. Then, there is a hybrid model based on the density profile from the Rosat all-sky survey. This last model is the standard one adopted for the interpretation.

To explore the range of predictions, these are also compared to the best-fitting polytropic model (Henriksen and Mushotzky 1986), which is clearly ruled out by subsequent observations. Table 7.1 lists the parameters used in these models.

TABLE 7.1
X-RAY MODEL PARAMETERS

Model	$n_{e,0}$ cm ⁻³	r_{core}	β	$T_{e,0}$ keV	T_{iso} keV	r_{iso}	γ	Remarks	Ref
Polytrope	3.5×10^{-3}	10.5	0.76	17.5	1.5		1
Hybrid N	3×10^{-3}	7.6	0.63	...	9.1	23	1.555	North Field	2
Hybrid A	3×10^{-3}	9.8	0.76	...	9.1	23	1.555	All IPC Fields	2
Rosat	2.89×10^{-3}	10.5	0.75	...	9.1	23	1.555	Hybrid Temperature Rosat	3

References: (1) Henriksen and Mushotzky 1986; (2) Hughes *et al.* 1988b; (3) Briel *et al.* 1992.

The top panel of figure 7.1 shows the radial behavior of the pencil-beam zero-frequency SZ effect $\Delta T_{0,true}(\theta)$ in the various models, while the bottom panel shows the effective SZ decrement $\Delta T_{0,eff}(\theta)$ seen by the 5.5m telescope configuration, given those same models. It can be seen that the polytropic model differs wildly from the others at small θ —this is due to the high gas temperature in the cluster’s center, which has a tremendous impact on the SZ effect while changing the X-ray emission only slightly. At larger radii, the polytropic model gives very similar predictions to the other models. The other models agree very well at almost all radii; this is no surprise since they are substantially similar with slightly different parameters.

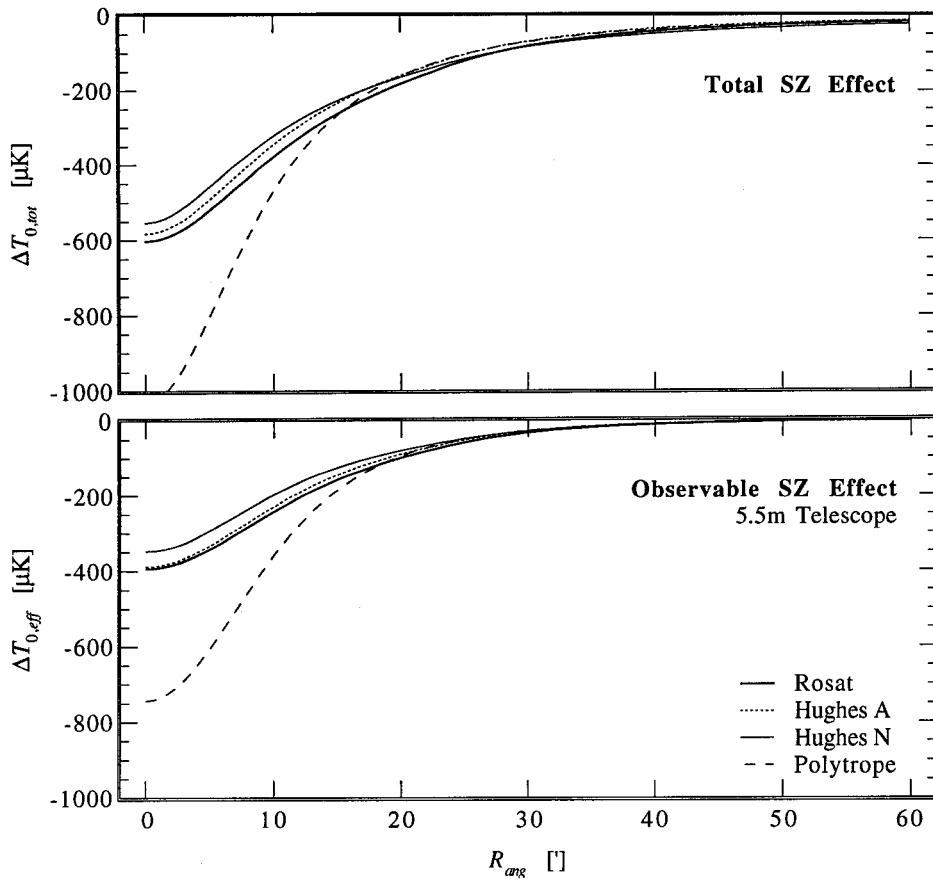


Figure 7.1 Predictions for the size of the SZ effect in the Coma cluster from various X-ray models (see text). The top panel shows the total strength of the SZ effect at an angular distance R from the cluster center, as if it was observed with a single pencil beam. The bottom panel shows the effective strength of the effect given the beam size and the beam separation of the 5.5m telescope.

Given the formal errors of the Rosat model (1.4% for the electron density, 7.7% for the electron temperature, 5.7% for the core radius, etc.), the formal statistical error of the X-ray model predictions is estimated to be about 10%. Comparing the different hybrid models in figure 7.1, this error range appears reasonable. The systematic error inherent in the choice of an X-ray model whose quality is judged on the basis of its X-ray rather than its SZ fit, is much more difficult. This question can be addressed best with high-quality X-ray temperature measurements by the ASCA telescope in the near future.

Ellipticity

To assess the influence of possible ellipticity of the cluster atmosphere, we simply took the hybrid Rosat model and elongated it along one axis in the plane of the sky. This produced a more strongly parallactic angle-dependent SZ prediction, but the overall value varied within a range of 5%. Clearly, the impact of the choice of the model is more important than possible ellipticity, which was therefore neglected.

7.2.3. Comparison with Observations—the Hubble Constant

Because of the quality of the all-sky Rosat data (Briel *et al.* 1992) and the constraints provided by the Tenma data (Hughes *et al.* 1988a), the hybrid Rosat model was chosen for a comparison with the observed SZ data. Before this is possible, we should explore the parallactic angle dependence of the predictions. This is shown in figure 7.2 for the actual SZ observing positions C, N, and CS.

It is clear from these figures that the parallactic angle variation in the central position C is very small and is only due to the small offset of the position to the southeast of the cluster center. As mentioned above, possible ellipticity in the cluster atmosphere increases this undulation, but only to a small fraction of the total value. Therefore, intrinsic parallactic angle structure based on the X-ray models does not explain the parallactic angle variation of the SZ results as described in the section on data quality above.

The parallactic angle variations of position N are far more dramatic, albeit at very small values. The smallness of the effect is due to the strongly reduced curvature in the SZ effect at large radii, when the line between the beams is perpendicular to the radius (zero parallactic angle). The positive value of the prediction at the parallactic angle extremes is due to the negative curvature, where the reference beam closest to the cluster center sees a much stronger SZ effect than the other two beams. Comparing this to the SZ result, we find some slight but marginal disagreement, given the statistical quality of the N data.

The parallactic angle predictions for the CS position are sedate and very close to zero. The observed value agrees with it completely. It is clear, therefore, that we do not detect a peculiar velocity effect from this presumably infalling subcluster.




Figure 7.2 Predicted effective SZ effect seen by the 5.5m telescope as a function of parallactic angle ψ , using the hybrid Rosat model (see text). The three panels show the prediction for the three observed positions in the cluster, C, N, and CS. The model was calculated with $H_0 = 50$ km/s/Mpc.

For position C, the only position from which a measurement of the Hubble constant can be made, the model predicts an observable value of $\Delta T_{0,m} = -368.1 \pm 37.0 \mu K$, roughly weighted by the integration time at each parallactic angle. Combined with the observed value of $\Delta T_{0,obs} = -302.1 \pm 48.0 \mu K$, this implies a Hubble constant of $H_{0,obs} = 74.2 \pm 27.9 \text{ km s}^{-1} \text{ Mpc}^{-1}$. While this error includes the formal errors of the X-ray model, it disregards possible systematic errors from the choice of the cluster model.

8. DISCUSSION

8.1 SYSTEMATIC ASTRONOMICAL EFFECTS

The calculation of the Hubble constant in the previous chapter makes several important assumptions. As mentioned before, it assumes that the length through the cluster along the line of sight equals the diameter of the cluster in the plane of the sky. In addition, the method assumes that the hot intracluster gas is not clumped, which could result in a significant overestimate of the gas content from X-ray observations. An excellent discussion of these sources of systematic error is given in Birkinshaw *et al.* (1991).

Both sources of systematic error cannot be assessed at this stage, and only estimates of their influence can be made. Most certainly, the departure from a spherical cluster (prolateness) is the most difficult to treat, since no method has been devised yet to measure it. Even more seriously, it can serve as a selection effect for the choice of the cluster. This is because all surface measures, such as X-ray surface brightness, galaxy surface density counts, and the SZ decrement itself are enhanced with a cluster that is extended along the line of sight.

Fortunately, observations of nearby clusters are relatively unaffected by this selection bias. The Coma cluster is a prime candidate for SZ observations for many reasons, and was not selected because it exhibited high surface density and brightness. In view of its apparent regularity on the sky, it would be very surprising if Coma was found to be strongly prolate along the line of sight. While such prolateness could certainly affect our estimate of the Hubble constant from this experiment, I expect such a readjustment to be less than a factor of 2. If there was an extension behind the cluster similar to the south-west extension seen in the X-ray images, it would affect the derived Hubble constant by well under 5%, assuming the low density and temperature suggested by earlier X-ray observations.

Measurements of the SZ effect in a complete sample selected by their total X-ray fluxes offer a reasonable chance to reduce our susceptibility to this prolate selection bias. However, such a catalog is difficult to compile, and will most likely be produced from the all-sky survey data from the Rosat telescope. In the mean time, a somewhat more hetero-

geneous catalog is available, based mainly on Exosat data (Edge *et al.* 1990). Data on these clusters are elaborated in Edge and Stewart (1991a, b).

Similarly, the impact of clumpiness of the gas atmosphere is much reduced with a nearby cluster. Some aspects of the non-smooth appearance of the X-ray atmosphere have already been taken into account in the X-ray models. Both Hughes *et al.* (1988b) and Briel *et al.* (1992) exclude obvious X-ray point sources from their density fits—this lessens the effect of known clumping of the X-ray atmosphere around single galaxies and also reduces the effect of background objects. While such editing is possible in the modeling of more distant clusters, it will be much less complete and effective than in the case of Coma. In this regard, the pointed Rosat observations of White *et al.* (1993) have much potential.

8.2 PECULIAR VELOCITY

The peculiar velocity of a cluster can also produce a secondary SZ effect, as explained in section 2.2. While this holds enormous promise in the long term to measure the peculiar velocities of clusters at arbitrary redshifts, its impact in the present experiment is that of a systematic error. However, given reasonable assumptions about Coma's peculiar velocity, we can find limits to the size of this effect.

The peculiar velocity effect follows the angular shape of the thermal effect very closely, especially at small radii, where the cluster atmosphere's temperature profile is flat. Therefore, we would expect to see a fraction of the peculiar velocity effect similar to the 61.0% of the thermal effect which we see with our observing configuration.

Using the value of the Hubble constant derived from our measurements, we can calculate the peak central optical depth of the X-ray cluster gas to be $\tau = 0.00463$. This makes the peculiar velocity effect small; if Coma's peculiar velocity was as large as $v_{pec} \approx 600 \text{ km s}^{-1}$ (see *e.g.*, Lauer and Postman 1993 for the magnitude of expected peculiar velocities), the observed SZ decrement would be corrupted by only about $15 \mu\text{K}$. This translates into an error in the derived Hubble constant of $\approx 7 \text{ km s}^{-1} \text{ Mpc}^{-1}$. Certainly, this error is small and can be disregarded at this stage.

8.3 ERROR BUDGET

The contributions to the total error in the observed effective SZ decrement of $\Delta T_o = -302.1 \pm 48 \mu\text{K}$ come from the following sources. The statistical error in the data of $29.2 \mu\text{K}$ and the allowance for undetected discrete sources of $35 \mu\text{K}$ (see section 6.4) share the bulk of the error. A small $13.9 \mu\text{K}$ contribution comes from the calibration of the telescope, and a $5.7 \mu\text{K}$ contribution comes from the subtraction of known discrete sources.

For the Hubble constant, the main part of the error comes from the total error in the SZ decrement and contributes $23.6 \text{ km s}^{-1} \text{ Mpc}^{-1}$. The formal uncertainty in the X-ray model contributes another $14.9 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

Once the source correction is improved so that the error allowance for undetected sources is substantially smaller, the total error estimate in the Hubble constant will be reduced to $22.0 \text{ km s}^{-1} \text{ Mpc}^{-1}$ in the absence of other improvements. On the other hand, if the X-ray model alone gained an accuracy of 3%, the total error in the Hubble constant would become $24.0 \text{ km s}^{-1} \text{ Mpc}^{-1}$. If both the undetected source allowance as well as the X-ray uncertainty could be reduced, the error would become $16.8 \text{ km s}^{-1} \text{ Mpc}^{-1}$. A further reduction of the error to $15.4 \text{ km s}^{-1} \text{ Mpc}^{-1}$ could be achieved by reducing the total error in the observed SZ decrement to 10%.

Of course, none of these error estimates can take into account the uncertainty inherent in the choice of the X-ray model. The uncertainty of this choice should be much reduced with further X-ray data.

9. CONCLUSIONS

This thesis describes the measurement of the Sunyaev-Zeldovich effect in the Coma cluster of galaxies. It is the first detection of the effect in a nearby cluster and is the first scientific result from the 5.5m telescope at the Owens Valley Radio Observatory.

Including all corrections for discrete sources and for the calibration of the telescope, the observed SZ effect for our observing configuration in the Coma cluster is $\Delta T_o = -302 \pm 48 \mu K$, which includes an error allowance for undetected discrete sources. This number is corrected for the frequency dependence of the SZ effect.

The interpretation of this result is made on the basis of existing high-quality X-ray models of the cluster atmosphere. In this context, the ratio of the effective value visible with our observing configuration to the total peak value of the effect when observed in the cluster center is 61.0%; from this we can infer a peak central SZ effect in the Coma cluster of $\Delta T_{peak} = -495.3 \pm 78.7 \mu K$.

Using the same interpretation and the black-body temperature of the cosmic microwave background of 2.726 ± 0.010 K (Mather *et al.* 1993), we can derive the peak central Comptonization to be $y_c = (9.06 \pm 1.44) \times 10^{-5}$. Similarly, the peak Thomson optical depth of the X-ray gas through the cluster center is $\tau_T = (4.63 \pm 0.74) \times 10^{-3}$.

From this observational result and its interpretation in the context of existing X-ray models, we can also derive a value for the Hubble constant of $H_0 = 74.2 \pm 27.9$ km s⁻¹ Mpc⁻¹.

This result will see many improvements in the coming years. First, a better and more complete subtraction of discrete radio sources will reduce the error of the observed SZ decrement significantly. In the meantime, of course, further observations will improve our statistical estimate of the data points obtained so far, but it will also allow for the addition of further points in the profile of the SZ effect across the cluster. This will enable us to form a self-consistent combination with an X-ray derived profile to assess our errors more realistically. Lastly, the quality of X-ray models of the Coma cluster will be improved dramatically by the inclusion of the pointed Rosat data (White *et al.* 1993) and dramatically improved temperature data from the ASCA satellite. A proposal to investigate Coma's temperature structure has been submitted with other collaborators.

Eventually, an entire map of the SZ effect in the cluster is possible. This would allow us to judge the effect of substructure in the cluster's X-ray atmosphere, but it would also permit an application of SZ observations to the study of this atmosphere to large radii.

Other nearby clusters can also be observed with the 5.5m telescope. Some candidates have already been observed in the summer of 1993, and more concentrated observations will certainly follow.

Observations of the SZ effect at higher frequency are interesting because of the lessened influence of discrete point sources and the eventual possibility to measure the peculiar velocity effect. This avenue is rapidly becoming possible with new high-frequency telescopes and existing experiments (*e.g.*, Wilbanks *et al.* 1993). Considerable effort should be spent to make this aspect of SZ observations as reliable as possible.

Finally, to reduce the impact of the prolate selection effect on our measurement of the Hubble constant from SZ measurements, a complete sample should be defined that is selected on the basis total X-ray flux. As mentioned in section 8.1, such samples are becoming available. It is only now, as measurements of the SZ effect are becoming routine observations, that SZ programs can begin to accommodate them.

REFERENCES

- Abell, G. O., Corwin, H. G., and Olowin, R. P. 1989, *Ap. J. Suppl. Series*, **70**, 1.
- Abramopoulos, F., Chanan, G. A., and Ku, W. H.-M. 1981, *Ap. J.*, **248**, 429.
- Arnaud, M., Hughes, J. P., Forman, W., Jones, C., Lachieze-Rey, M. *et al.* 1992, *Ap. J.*, **390**, 345.
- Baars, J. W. M., Genzel, R., Pauliny-Toth, I. I. K., and Witzel, A. 1977, *Astron. Astrophys.*, **61**, 99.
- Bahcall, J. N. and Sarazin, C. L. 1978, *Ap. J.*, **219**, 781.
- Becker, R. H., White, R. L., and Edwards, A. L. 1991, *Ap. J. Suppl. Series*, **75**, 1.
- Birkinshaw, M., Gull, S. F., and Hardebeck, H. 1984, *Nature*, **309**, 34.
- Birkinshaw, M., Gull, S. F., and Northover, K. J. E. 1981, *M. N. R. A. S.*, **197**, 571.
- Birkinshaw, M., Gull, S. F., Hardebeck, H. E., and Moffet, A. T. 1993, *Ap. J.*, submitted.
- Birkinshaw, M., Hughes, J. P., and Arnaud, K. A. 1991, *Ap. J.*, **379**, 466.
- Bond, J. R. and Myers, S. T. 1993, *Ap. J.*, submitted.
- Briel, U. G., Henry, J. P., and Böhringer, H. 1992, *Astron. Astrophys.*, **259**, L31.
- Cavaliere, A. and Fusco-Femiano, R. 1976, *Astron. Astrophys.*, **49**, 137.
- Cavaliere, A. and Fusco-Femiano, R. 1978, *Astron. Astrophys.*, **70**, 677.
- Condon, J. J., Broderick, J. J., and Seielstad, G. A. 1989, *A. J.*, **97**, 1064.
- Dressler, A. and Gunn, J. E. 1992, *Ap. J. Suppl. Series*, **78**, 1.
- Edge, A. 1989, Ph.D. Thesis, Leicester University.
- Edge, A. C. and Stewart, G. C. 1991a, *M. N. R. A. S.*, **252**, 414.
- Edge, A. C. and Stewart, G. C. 1991b, *M. N. R. A. S.*, **252**, 428.
- Edge, A. C., Stewart, G. C., Fabian, A. C., and Arnaud, K. A. 1990, *M. N. R. A. S.*, **245**, 559.
- Fitchett, M. J. 1988, in *Clusters of Galaxies and Large-Scale Structure*, p. 143; ed. J. M. Dickey (San Francisco: Astronomical Society of the Pacific)
- Gaier, T., Schuster, J., Gundersen, J., Koch, T., Seiffert, M., Meinhold, P. R., and Lubin, P. M. 1992, *Ap. J.*, **398**, L1.
- Gull, S. F. and Northover, K. J. E. 1976, *Nature*, **263**, 572.
- Handbook of Chemistry and Physics (HCP), Chemical Rubber Company.
- Henriksen, M. J. and Mushotzky, R. F. 1986, *Ap. J.*, **302**, 287.
- Herbig, T. and Readhead, A. C. S. 1991, *OVRO Bandpass Memorandum*.
- Herbig, T., Birkinshaw, M., and Arnaud, M., 1993, in preparation.
- Herbig, T., Lawrence, C. R., Readhead, A. C. S., and Gulkis, S. 1993, *Nature*, submitted (*cf.* appendix A).
- Hogan, C. J. 1992, *Ap. J.*, **398**, L77.
- Hughes, J. P. 1989, *Ap. J.*, **337**, 21.
- Hughes, J. P., Gorenstein, P., and Fabricant, D. 1988, *Ap. J.*, **329**, 82.
- Janssen, M. A., Golden, L. M., and Welch, W. J. 1974, *Astron. Astrophys.*, **33**, 373.

- Jones, M., Saunders, R., Alexander, P., Birkinshaw, M., Dillon, N., *et al.* 1993, *Nature*, in press.
- Klein, M. J. and Gulkis, S. 1978, *Icarus*, **35**, 44.
- Kompaneetz, A. S. 1957, *Sov. Phys. JETP*, **4**, 730.
- Kraus, J. D. 1982, *Radio Astronomy* (Ohio: Cygnus-Quasar).
- Lake, G. and Partridge, R. B. 1977, *Nature*, **270**, 502.
- Lake, G. and Partridge, R. B. 1980, *Ap. J.*, **237**, 378.
- Lawrence, C. R., Herbig, T., and Readhead, A. C. S. 1993, *Proc. of the IEEE*, in press (*cf.* appendix B).
- Leir, A. A. and van den Bergh, S. 1977, *Ap. J. Suppl. Series*, **34**, 389.
- Markevich, M., Blumenthal, G. R., Forman, W., Jones, C., and Sunyaev, R. A. 1991, *Ap. J.*, **378**, L33.
- Markevich, M., Blumenthal, G. R., Forman, W., Jones, C., and Sunyaev, R. A. 1992, *Ap. J.*, **395**, 326.
- Mather, J. C., Cheng, E. S., Cottingham, D. A., Eplee, R. E., Fixsen, D. J., *et al.* 1993, submitted to *Ap. J.*
- Myers, S. T., Readhead, A. C. S., and Lawrence, C. R. 1993, *Ap. J.*, **405**, 8.
- NBS Circular 514, *Table of Dielectric Constants of Pure Liquids*, National Bureau of Standards, p. 3.
- Nyquist, H. 1928, *Phys. Rev.*, **32**, 110.
- Parijsky, Yu. N. 1972, *Astronomicheskij Zhurnal*, **49**, 1322.
- Pearson, T. J. 1993. *OVRO 40m and 5.5m Antennas: Control System*.
- Perrenod, S. C. and Lada, C. J. 1979, *Ap. J.*, **234**, L173.
- Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T. 1986, *Numerical Recipes* (Cambridge: Cambridge University Press)
- Readhead, A. C. S., Lawrence, C. R., Myers, S. T., Sargent, W. L. W., Hardebeck, H. E., and Moffet, A. T. 1989, *Ap. J.*, **346**, 566.
- Rudnick, L. 1978, *Ap. J.*, **223**, 37.
- Rybicki, G. B. and Lightman, A. P. 1979, *Radiative Processes in Astrophysics* (New York: Wiley)
- Sarazin, C. L., Rood, J. H., and Struble, M. F. 1982, *Astron. Astrophys.*, **108**, L7.
- Sargent, W. L. W. 1973, *Pub. Astron. Soc. Pac.*, **85**, 281.
- Schallwisch, D. and Wielebinski, R. 1979, *Astron. Astrophys.*, **71**, L15.
- Silk, J. and White, S. D. M. 1978, *Ap. J. Letters*, **226**, L103.
- Silverberg, R., Cheng, E., Cottingham, D., Fixsen, D., Gentieu, P., *et al.* 1993, *Bull. Amer. Astron. Soc.*, **24** (4), 1264.
- Smoot, G. F., Bennett, C. L., Kogut, A., Wright, E. L., Aymon, J. *et al.* 1992, *Ap. J.*, **396**, L1.
- Steele, G. R. 1990, *Common LISP: The Language*, Second Ed. (Digital Press).
- Struble, M. F. and Rood, H. J. 1987, *Ap. J. Suppl. Series*, **63**, 543.
- Sunyaev, R. A. and Zeldovich, Ya. B. 1980a, *M. N. R. A. S.*, **190**, 413.
- Sunyaev, R. A. and Zeldovich, Ya. B. 1980b, *Ann. Rev. Astron. Astrophys.*, **18**, 537.
- Tarter, J. C. 1978, *Ap. J.*, **220**, 749.
- Thompson, A. R. and D'Addario, L. R. 1982, *Radio Science*, **17**, 357.
- Uson, J. 1987, in *Radio Continuum Processes in Clusters of Galaxies*, NRAO Green Bank Workshop **16**, 255; eds. O'Dea, C and Uson, J., (Green Bank, WV: NRAO).
- Venturi, T. 1985, Ph.D. thesis, University of Bologna.

- Venturi, T., Giovannini, G., and Feretti, L. 1990, *A. J.*, **99**, 1381.
- White, S. D. M., Briel U. G., and Henry, J. P. 1993, *M. N. R. A. S.*, **261**, L8.
- Wilbanks, T. M., Ade, P. A. R., Fischer, M. L., Ho, T. R., Holzapfel, W. L. *et al.* 1993, *Bull. Amer. Astron. Soc.*, **25** (1), 740.
- Willson, M. A. G 1970, *M. N. R. A. S.*, **151**, 1.
- Zeldovich, Ya. B. and Sunyaev, R. A. 1969, *Astrophys. Space Sci.*, **4**, 301.

APPENDIX A

The First Detection of the Sunyaev-Zeldovich Effect in a Nearby Cluster of Galaxies

T. HERBIG,[†] C. R. LAWRENCE, & A. C. S. READHEAD
Owens Valley Radio Observatory, California Institute of Technology

S. GULKIS
Jet Propulsion Laboratory

We report the detection of the Sunyaev-Zeldovich effect toward the Coma cluster of galaxies at a significance level of 9σ , using the new 5.5m telescope of the Owens Valley Radio Observatory. After correcting for discrete radio sources and interpreting the measurement in the context of existing X-ray models, we derive a value of $\Delta T_{0,C} = -495 \pm 79 \mu\text{K}$ at the cluster center. This is the first detection of the SZ effect in a nearby cluster—a crucial development since the usefulness of SZ observations generally relies upon comparisons with optical and X-ray observations, which are much more detailed for nearby clusters. Comparison with existing X-ray models yields a Hubble constant of $74 \pm 28 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

The Sunyaev-Zeldovich (SZ) effect describes the spectral distortion of the cosmic microwave background radiation due to inverse Compton cooling of the X-ray gas in a cluster of galaxies¹. In the Rayleigh-Jeans regime this leads to a reduction in the brightness temperature of the background. Combined with X-ray data, the SZ effect provides a powerful probe of the intracluster medium and a direct measurement of the Hubble constant^{2,3}.

Since it was detected reliably nine years ago⁴, the SZ effect has been measured with various techniques in four distant clusters of galaxies^{5,6,7,8} with redshifts ranging from 0.17 to 0.55^{9,10}.

Despite the distance independence of the effect, it has heretofore been observed only in distant clusters, since the resolution of large single dishes and interferometers now in use is too high to permit observations of nearby clusters with large angular diameters. Even in distant clusters, the SZ effect is observable on only a few of the shortest baselines of the Ryle telescope⁸ or of the Caltech mm-array. This is unfortunate, since the full astronomical and cosmological potential of SZ observations hinges on their combination with high-quality X-ray data, which are more difficult to obtain for distant clusters.

It should be clear, therefore, that single-dish and interferometric observations of the SZ effect generally complement each other, and that the combination of observations using *both* techniques will provide the best line of attack in the future. It should also be clear that single dish observations with a small telescope have a particularly vital role to play in nearby clusters.

The Coma cluster, Abell 1656, at a redshift of 0.0235¹¹, is a prime candidate for such observations. One of the best-studied clusters, it is comparatively regular¹², has an Abell richness of 2¹³, and contains a hot X-ray atmosphere¹⁴ that should produce a strong SZ effect.

[†] Also at Physics Department, Princeton University

HERBIG ET AL.

During the winters of 1992 and 1993, we observed the Coma cluster with the new 5.5m diameter radio telescope at the Owens Valley Radio Observatory. This telescope is dedicated to cosmic microwave background observations. Its receiver operates at a center frequency of 32 GHz with a bandwidth of 5.7 GHz. The input of the cryogenic HEMT amplifier is switched every millisecond between two 7' beams separated by 22'2 on the sky in azimuth. The millisecond-sampling backend is fully digital to reduce potential systematic effects and to permit thorough testing of the entire system.

During good observing weather in the 1993 season, the total system temperature at the zenith was about 54 K, including the cosmic microwave background, atmospheric emission of 10 K, and total ground pickup of 9 K¹⁵. On the basis of the flux density of DR21¹⁶, the main beam efficiency was measured to be 74%. The low point-source gain of 4.56 mK Jy⁻¹ is a distinct advantage in reducing discrete source contamination.

The 5.5m telescope has a specially designed feed leg structure to minimize differential ground pickup¹⁵. As described below, we perform double-switched flux measurements to reduce systematic effects such as ground spillover. Nevertheless, the residual differential ground pickup reaches 150 μ K at some zenith angles.

We therefore adopted a triple differencing technique. The first two levels follow a standard double switching scheme¹⁷ with Dicke switching between the two beams on the sky at 500 Hz and position switching every 20s to 30s, producing flux measurements with typical durations of 100s. The third level is aimed at reducing differential ground pickup and involves alternating between the target field and a blank leading or trailing reference field 15 minutes away in time, so that the telescope tracks the two fields over precisely the same path relative to the ground. This produces pairs of double-switched flux measurements that were taken in very similar atmospheric and ground conditions.

With this approach, we observed two points in the cluster—a central position (C), offset by 3'5 to the SE of the X-ray center to minimize discrete source contamination, and an off-center position 30' to the north of the cluster center (N), where we expect to measure a negligible SZ decrement. In addition, a position (CS) centered on the south-western extension of the X-ray atmosphere¹¹ was observed in 1993.

During both seasons, the observations started in the middle of February and continued until May, with Coma overhead at night. This was a significant advantage for our observations, since night-time data are more stable than those during the day. The weather during both observing seasons was poor—more than half of the data had to be rejected, especially in 1993. Atmospheric conditions deteriorated noticeably in late spring, when data on positions N and CS were gathered.

To ensure proper ground subtraction, data pairs with a position discrepancy of more than 2' with respect to the ground were eliminated. To identify data taken during bad weather, we first discarded all individual flux measurements whose internal scatter estimate exceeded the expected thermal noise by a factor of 2. We further identified all groups of these high-noise data and removed all preceding and following neighbors within a margin equal to the duration of the rejected group (up to a maximum of 2 hours). This minimizes our susceptibility to weather-induced biases not manifested by short-term fluctuations. In all editing procedures, corresponding flux measurements were removed pairwise to retain proper ground subtraction. These editing methods are robust: the exact choice of editing parameters made little difference in the final result.

The edited overall results from the observations are given in Table 1, showing Rayleigh-

SZ EFFECT IN COMA

Jeans sky brightness temperatures corrected to zero frequency for the spectral dependence of the SZ effect. Apart from the differencing inherent in the observing technique, no offsets or drifts were subtracted from these values. The cumulative result for the central position is a 9.1σ detection of the SZ effect—the result is highly significant in both observing seasons and in both leading and trailing subsets independently. The cumulative results for positions N and CS are consistent with zero; the data are of poorer quality than those for position C due to shorter total integration times and worse atmospheric conditions. Among other aspects of the data, this null result gives us confidence that our observing technique cancels systematic effects such as residual ground pickup.

TABLE 1
Observed Values

Position	Season	$\Delta T_{0,obs}$	N_{pairs}
C	1992	$-261.9 \pm 33.2 \mu\text{K}$	1518
	1993	-269.8 ± 60.8	955
	overall	-263.7 ± 29.2	2473
N	1992	-47.0 ± 64.0	583
	1993	-128.9 ± 81.7	295
	overall	-78.2 ± 50.4	878
CS	1993	-17.6 ± 71.2	291

The observed values must be corrected for contamination by discrete radio sources. A preliminary list was identified from the 5C4 and Condon surveys^{18,19} and observed with the OVRO 40m telescope at 18.5 GHz to obtain spectral information. Only two of these sources contribute significantly to the observed SZ effect: 5C4.81 and 5C4.85, as well as Coma C²⁰, the cluster's central halo source. The correction for their influence on the center beam amounts to a total of $38.4 \pm 5.7 \mu\text{K}$. To account for sources with flatter spectra ($\alpha > -0.51$) that fall below the flux density limits of those two surveys, an additional allowance of $\pm 35 \mu\text{K}$ is made. A sensitive VLA survey is under way to eliminate this allowance. The corrected observed SZ decrement in position C thus becomes $\Delta T_0 = -302.1 \pm 48.0 \mu\text{K}$.

ΔT_0 is derived with the assumption that the cluster has circular symmetry on the sky. If this should not be the case, then the observed decrement will depend on the parallactic angle of the azimuthally separated reference beams, and a single value will not be sufficient to describe the cluster. Our present data do not show statistically significant variations of the decrement with parallactic angle. However, this question can be settled only with a true image of the effect, which currently can be obtained only by multiple pointings of the 5.5m telescope due to the large angular size of the cluster.

The result in the center position can be interpreted in the context of existing X-ray models^{14,21} of the gas atmosphere, which predict that our observing configuration measures 61.0% of the central peak SZ effect at position C. From this, we derive a central SZ decrement in Coma of $\Delta T_{0,C} = -495 \pm 79 \mu\text{K}$. The increased error bar reflects the uncertainty allowance of the discrete source subtraction, and is not indicative of the statistical reliability of the detection. The peak Comptonization is thus $y_C = (9.06 \pm 1.44) \times 10^{-5}$ and the peak Thomson optical depth is $\tau_T = (4.63 \pm 0.74) \times 10^{-3}$.

Comparison of our observations with these X-ray models yields a Hubble constant of $74 \pm 28 \text{ km s}^{-1} \text{ Mpc}^{-1}$. The error in this value reflects the statistical uncertainty of the SZ observations, the calibration of the system, the error in the corrections for known discrete radio sources, the allowance for undetected discrete sources, and the quoted uncertainty of

HERBIG ET AL.

the X-ray model. However, it does not allow for the systematic error that originates from the choice of the X-ray model. The possibility that the cluster atmosphere may not be spherically symmetric or that it may be clumped on scales less than the X-ray resolution³ could affect this result, and should be addressed in the future. Similarly, peculiar motion of the cluster of up to 600 km s^{-1} would introduce a small additional Doppler-based SZ effect²², affecting our value of the Hubble constant by $7 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

Given the dramatic progress in SZ measurements over the past few years, both with single dishes and with interferometers, the major uncertainty in their interpretation now lies in the choice of the X-ray model of the cluster atmosphere, which will be refined significantly by present and future observations with the Rosat and ASCA telescopes.

This project would have been impossible without the dedication of our engineers at OVRO, H. Hardebeck, M. Hodges, and R. Keeney. Further thanks go to M. Birkinshaw, S. T. Myers, and J. Baker for many fruitful discussions. We appreciate the efforts of the National Radio Astronomy Observatory, TRW, and the Jet Propulsion Laboratory in designing and building the HEMT devices and amplifiers critical for this project. Microwave background studies at OVRO are supported under a grant from the National Science Foundation.

REFERENCES

- [1] Zeldovich, Ya. B. and Sunyaev, R. A. 1969, *Astrophys. Space Sci.*, **4**, 301-316.
- [2] Silk, J. and White, S. D. M. 1978, *ApJ Letters*, **226**, L103-106.
- [3] Birkinshaw, M., Hughes, J. P., and Arnaud, K. A. 1991, *ApJ*, **370**, 466-481.
- [4] Birkinshaw, M., Gull, S. F., and Hardebeck, H. 1984, *Nature*, **309**, 34-35.
- [5] Birkinshaw, M., Gull, S. F., Hardebeck, H. E., and Moffet, A. T. 1993, *ApJ*, submitted.
- [6] Uson, J. 1987, in *Radio Continuum Processes in Clusters of Galaxies*, NRAO Green Bank Workshop **16**, 255; eds. O'Dea, C and Uson, J., NRAO Green Bank, WV.
- [7] Wilbanks, T. M., Ade, P. A. R., Fischer, M. L., Ho, T. R., Holzapfel, W. L., et al. 1993, *Bull. Amer. Astron. Soc.*, **16** (1), 740.
- [8] Jones, M., Saunders, R., Alexander, P., Birkinshaw, M., Dillon, N. et al. 1993, *Nature*, in press.
- [9] Struble, M. F. and Rood, H. J. 1987, *ApJ Suppl Series*, **63**, 543-553.
- [10] Dressler, A. and Gunn, J. E. 1992, *ApJ Suppl Series*, **78**, 1-60.
- [11] Sarazin, C. L., Rood, J. H., and Struble, M. F. 1982, *Astron. Astrophys.*, **108**, L7-10.
- [12] White, S. D. M., Briel U. G., and Henry, J. P. 1993, *Mon. Not. R. A. Soc.*, **261**, L8-L12.
- [13] Abell, G. O., Corwin, H. G., and Olowin, R. P. 1989, *ApJ Suppl Series*, **70**, 1-138.
- [14] Hughes, J. P., Gorenstein, P., and Fabricant, D. 1988, *ApJ*, **329**, 82-96.
- [15] Lawrence, C. R., Herbig, T., and Readhead, A. C. S. 1993, *Proc. of the IEEE*, in press.
- [16] Klein, M. J. and Gulkis, S. 1978, *Icarus*, **35**, 44.
- [17] Readhead, A. C. S., Lawrence, C. R., Myers, S. T., Sargent, W. L. W., Hardebeck, H. E., and Moffet, A. T. 1989, *ApJ*, **346**, 566-587.
- [18] Willson, M. A. G. 1970, *Mon. Not. R. A. Soc.*, **151**, 1-44.
- [19] Becker, R. H., White, R. L., and Edwards, A. L. 1991, *ApJ Suppl Ser*, **75**, 1-229.
- [20] Venturi, T., Giovannini, G., and Feretti, L. 1990, *Astron. J.*, **99**, 1381-1396.
- [21] Briel, U. G., Henry, J. P., and Böhringer, H. 1992, *Astron. Astrophys.*, **259**, L31-L34.
- [22] Sunyaev, R. A. and Zeldovich, Ya. B. 1980, *Mon. Not. R. A. Soc.*, **190**, 413-420.

APPENDIX B

REDUCTION OF GROUND SPILLOVER IN THE OWENS VALLEY 5.5 m TELESCOPE

C. R. LAWRENCE, T. HERBIG, AND A. C. S. READHEAD

Owens Valley Radio Observatory, Mailcode 105-24
California Institute of Technology, Pasadena, CA 91125

ABSTRACT

The 5.5m telescope of the Owens Valley Radio Observatory is used for highly sensitive measurements of the cosmic microwave background radiation, which are easily contaminated by ground spillover. We describe the modifications made to the secondary mirror support structure that reduced the total ground pickup of the antenna from 27 K to 9 K and decreased the peak differential ground pickup for our observing technique from 4.3 mK to 140 μ K. This was achieved by reducing the physical obstruction, arranging the support leg angles to minimize direct reflections, and installing radiation baffles to control the direction of the reflection or scatter lobes.

1. INTRODUCTION

The 5.5 m telescope at the Owens Valley Radio Observatory (OVRO) is dedicated to sensitive observations of the cosmic microwave background radiation. We measure the second derivative of sky temperature over a 44' ($0^\circ 74'$) scale with a double switching technique that combines Dicke switching between two feeds with position switching of the telescope (see e.g. [1]). This technique cancels all offsets and linear effects, leaving us sensitive only to curvature in the radiation field between the three positions. Systematic errors in these measurements must be less than 10 μ K.

The telescope, a Cassegrain system with an altitude-azimuth mount, has a Dicke-switched receiver with a center frequency of 32 GHz, a bandwidth of 6 GHz, and a receiver noise temperature of 23 K including the feeds. The receiver is sensitive to left circular polarization. The two corrugated scalar feeds, separated in azimuth by 22:2 ($0^\circ 37'$), produce a 20 dB taper at the edge of the primary dish. The beam size is 7' ($0^\circ 12'$). Because of the underillumination, the telescope has an aperture efficiency of 55%.

When the telescope first came into operation in December 1990, it became apparent that it suffered from unacceptably large ground pickup. There was a 27 ± 4 K contribution to the total system temperature, roughly independent of telescope orientation. More importantly for our observations, the difference in ground pickup between the central beam and the average of the beams 22:2 away (hereafter referred to as differential ground pickup) was as large as 4.3 mK in some positions of the telescope. This differential ground pickup was measured using our standard double-switched observing technique [1], except that instead of tracking a source across the sky, we observed positions fixed in azimuth and zenith angle. We were extremely pessimistic that such a large effect could be removed to our required accuracy of less than 10 μ K.

The source of the large ground pickup proved to be the support structure for the secondary reflector (see Figure 1). This structure appears to have been designed for maximum mechanical stability. The total geometrical blockage, weighted by the aperture illumination, was about 7%, while the double legs were at an angle of $\sim 47^\circ$ to the telescope axis. Thus the telescope had large sidelobes at 94° . Not only did this lead to large total ground pickup,

LAWRENCE, HERBIG, AND READHEAD

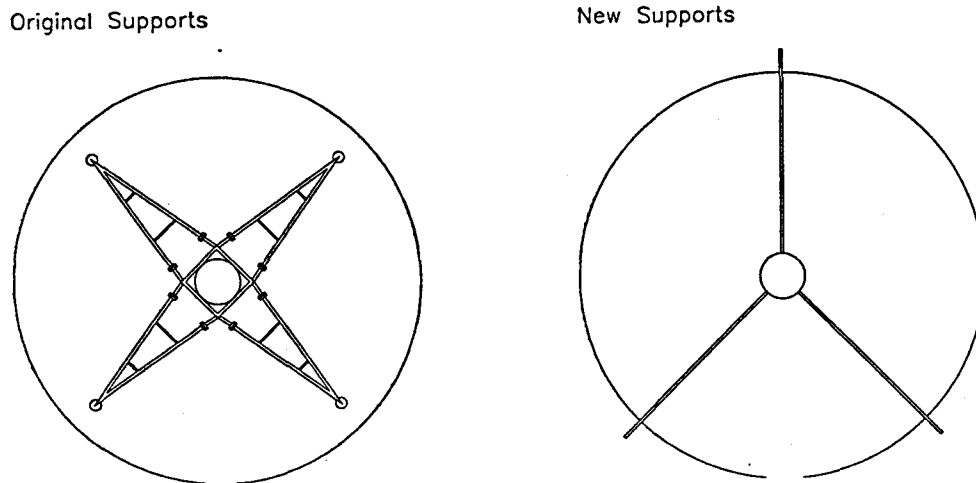


FIG 1.—Front view of the original support structure (left) and the new supports (right) as seen with the telescope pointed to the horizon. It can be seen that the total blockage was reduced and that the legs were led out over the edge of the dish. This view also shows the inverted Y configuration used in our system. The two lower legs are separated by 90° rather than 120° to align them with the backup structure of the 16 radial panels that form the dish.

but because the telescope is small and close to the ground, variations in emissivity among nearby objects led to unacceptably large differential signals.

The ideal solution to this problem, replacing the antenna with a clear aperture design [2] to achieve an extremely clean beam, was beyond our means. Instead, we replaced the secondary support structure, starting with a design used for several antennas of similar size at the Hat Creek Radio Observatory. Support strut scattering has been discussed widely in the literature (e.g. [3-5 and references therein]), but usually with more emphasis on the reduction of total scattered power than of differential scattering. Our new supports achieve low total geometrical blockage and good mechanical stability, and at the same time control reflections so that only sky radiation can enter the feeds in one reflection for telescope zenith angles up to 60° .

2. GEOMETRICAL BLOCKAGE

A significant reduction in geometrical obstruction was achieved by replacing the original double-leg quadrupod with a simple tripod (Figure 1), reducing the number of legs from 8 to 3. The new legs are trusses made of two $38\text{ mm} \times 76\text{ mm}$ aluminum beams, separated along

GROUND SPILLOVER

the telescope axis by cross bracing. Moreover, since the new legs extend beyond the edge of the dish, there is no blockage in the converging beam of the telescope. We also removed a strut cage circumscribing the secondary reflector, which was a significant obstruction in the center of the aperture illumination. Weighted by the illumination, these changes reduced the total physical obstruction by about a factor of four.

3. SUPPORT LEG GEOMETRY

Landecker *et al.* (1991) considered the effect of various aspects of support leg geometry on ground pickup in radio telescopes. We have extended their arguments and specialized them to our telescope, making use of the fact that the inner face of the support legs has a width of 4λ at our observing frequency. Therefore, diffraction effects have an opening angle of only $\sim 14^\circ$, and support leg reflections can be directed to the sky in a quasi-geometric fashion.

As shown in Figure 2, when the angle α between the support legs and the axis of the telescope is 60° the antenna itself prevents paraxial rays from reaching the ground in a *single* reflection off the legs. This is a significant change from the previous arrangement, where paraxial rays from much of the dish hit the ground close to the telescope at all zenith angles. It can be seen from the figure that some secondary reflections can still reach the ground at zenith angles of about 30° or larger. Because of the strong taper of the primary's illumination, these secondary reflections contain little power. Nevertheless, as we describe in the next section, the addition of inexpensive sheet metal baffles to the inner edges of the legs can further reduce these secondary reflections.

We also investigated the merits of an upright versus an inverted **Y** configuration of the tripod. For Landecker *et al.* the dominant consideration was that with an upright **Y** there is only one leg at the bottom whose scattering cone intersects the ground at large zenith angles. In the next section, however, we describe a way to largely eliminate that problem. We were then free to choose an inverted **Y** configuration, which has the practical advantage of allowing much simpler access to the front of the receiver with a ladder from the ground.

4. RADIATION BAFFLES

Landecker *et al.* also discussed the advantages of triangular support legs in reducing ground pickup. Scattering from a flat surface is more concentrated into a single direction than that from a cylindrical surface. That direction, of course, is determined by the orientation of the flat surface. The cross-section of the legs in the Landecker *et al.* telescope was fixed by structural requirements. For our small telescope, however, it was easy to add non-structural reflecting baffles that could be built to any tilt angle β on the inside surface of the support legs (Figure 3). By adjusting β , we can place all reflections directly on the sky for a wide range of telescope zenith angles.

Imagine a paraxial ray from the dish reflecting off the baffle. As shown in Figure 4, the direction of the ray after reflection can be described by two angles: φ , the angle between the incident and reflected rays (equivalently, the angle between the reflected ray and the telescope axis); and ρ , the angle between the projections of the support leg and the reflected ray in the plane of the dish. φ and ρ can be calculated from the support leg angle α and the baffle tilt angle β :

$$\sin \frac{\varphi}{2} = \cos \beta \sin \alpha \quad \text{and} \quad \tan \rho = \frac{\tan \beta}{\cos \alpha}.$$

LAWRENCE, HERBIG, AND READHEAD




FIG 2.—Side view of the reflecting surfaces and one of the new support legs at an inclination angle of $\alpha = 60^\circ$ to the telescope axis. With this arrangement, no paraxial rays from the dish can reach the ground in a single reflection. However, the rays on the periphery of the beam experience secondary reflections into large angles from the telescope axis.

These angles are plotted in Figures 5a and b as functions of β , for various values of α .

As $\rho \leq 90^\circ$, reflections from the upper leg are always directed downward. Since we must be able to observe sources to zenith angles of 60° with the local horizon (set by mountains) below $\gtrsim 85^\circ$ for most azimuths, we require $\varphi \lesssim 25^\circ$, giving $\beta \gtrsim 75^\circ$. We used $\beta = 75^\circ$ with a symmetrical arrangement as in Figure 3. For the lower support legs, however, we can use an asymmetrical design that directs the sidelobes up and away from the ground. Since these sidelobes should be well above the horizon when the telescope is pointed to the zenith, an angle of $\beta = 60^\circ$ was chosen. The resulting baffles direct the reflection sidelobes away from the ground at all zenith angles less than 60° .

5. RESULTS

The improvement obtained by changing the support structure is illustrated in Figure 6 and quantified in Table 1. Taken together, the new supports lowered the total ground pickup from 27 K to 9 K. We suspect that most of the remaining ground pickup is produced by elements other than the secondary support legs. Of much more importance for our observations, the maximum residual differential ground pickup was reduced by a factor of 30, from 4.3 mK to 140 μ K. In observations spanning a large, continuous range of zenith

GROUND SPILLOVER




FIG 3.—Cross section of the top support leg and one of the bottom legs showing the radiation baffle and its tilt angle β . The baffle for the top leg is symmetric, while those for the bottom legs have their tilted faces turned upward to direct reflections away from the ground.

angles, the residual differential ground pickup averages to about $90 \mu\text{K}$, consistent with the remainder seen in the test data from Figure 6. At some zenith angles, the improvement is even better. Near the North celestial pole, at zenith angles of 51° , the 1σ upper limit to residual differential ground pickup is $36 \mu\text{K}$. With the new supports, the main beam efficiency was measured to be 76%.

TABLE 1
TOTAL AND DIFFERENTIAL GROUND PICKUP

	Old Support Structure	New Support Structure
Total	$27 \pm 4 \text{ K}$	$9 \pm 1 \text{ K}$
Differential		
peak	$4.3 \pm 1.0 \text{ mK}$	$-0.14 \pm 0.04 \text{ mK}$
most significant	$-2.48 \pm 0.13 \text{ mK}$	$0.10 \pm 0.02 \text{ mK}$

As expected, the dominant improvement came from the change in overall geometry. The baffles should make the most difference at large zenith angles. Unfortunately, this is where increased atmospheric noise makes small effects particularly difficult to measure, and

LAWRENCE, HERBIG, AND READHEAD




FIG 4.—Schematic view of the telescope showing all three support legs and the angles ϕ and ρ that describe the direction of a ray after reflection from a radiation baffle on the leg's underside.

the data taken at large zenith angles before the baffles were installed were not sufficiently accurate to quantify the difference. Therefore, we have combined both sets for the data shown in Figure 6b.

Given the level of remaining differential ground pickup, we are confident that an observing strategy using proper ground subtraction techniques can achieve the level of systematic error that we require. From data observed over a large range of zenith angles [6] using ground subtraction, we find a 1σ upper limit of $43\ \mu\text{K}$ for offsets due to incorrect ground subtraction.

We would like to thank J. Welch for his advice and the design of the Hat Creek telescopes. We also thank T. Seling, D. Bathker, P. Cramer, R. Levy, and W. Schaal for their help in various aspects of this work.

REFERENCES

- [1] A. C. S. Readhead, C. R. Lawrence, S. T. Myers, W. L. W. Sargent, H. E. Hardebeck, and A. T. Moffet, "A limit on the anisotropy of the microwave background on arc minute scales," *Astrophysical Journal*, vol. 346, pp. 566-587, 1989.
- [2] A. G. Cha, "Design for a 1.5 m unblocked aperture antenna," *IEEE Trans Antennas & Prop AP-31*,

GROUND SPILLOVER

- vol. 6, pp. 896-902, 1983.
- [3] T. L. Landecker, M. D. Anderson, D. Routledge, R. J. Smegal, P. Trikha, and J. F. Vaneldik, "Ground radiation scattered from feed support struts: a significant source of noise in paraboloidal antennas," *Radio Science*, vol. 26 (2), pp. 363-373, 1991.
 - [4] H. Thielen, "Reduction of the Strut Radiation of Reflector Antennas," *IEEE/APS Symposium*, pp. 504-507, 1982.
 - [5] A. Hunt and A. Wright, "Complex, off-axis sidelobes of a radio telescope caused by feed-support legs," *Monthly Not. Royal Astron. Soc.*, vol. 258, pp. 217-224, 1992
 - [6] T. Herbig, C. R. Lawrence, A. C. S. Readhead, and S. Gulkis, "Detection of the SZ effect in the Coma cluster of galaxies", submitted to *Nature*




FIG 5.—Angle from the telescope axis φ (a) and the azimuthal angle ρ (b) as a function of the baffle tilt angle β . Shown are curves at various support leg inclination angles α .

LAWRENCE, HERBIG, AND READHEAD

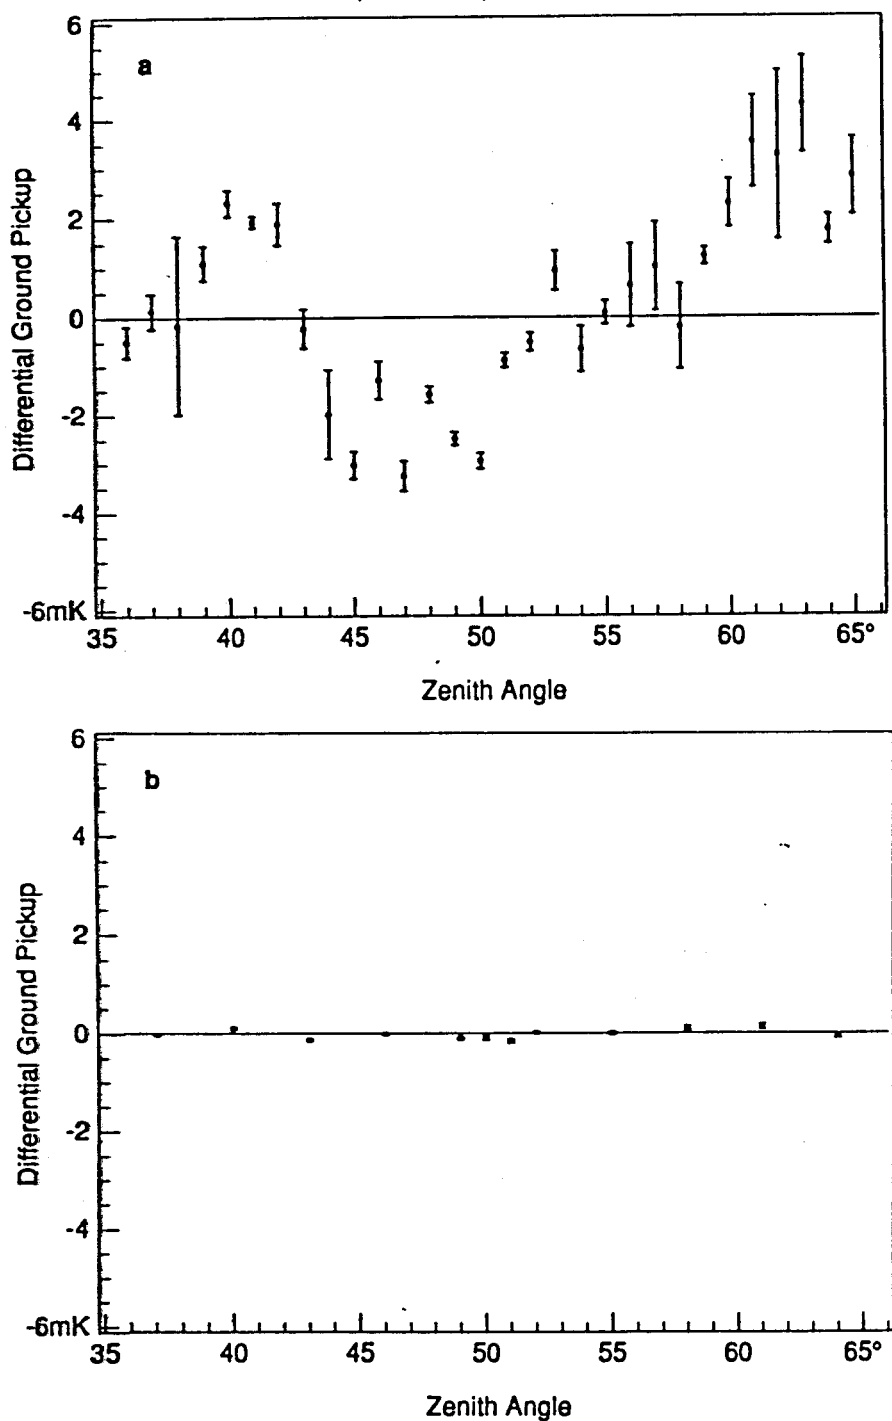


FIG 6.—Double-switched flux density measurements of blank sky at a series of fixed zenith angles, in mK of equivalent Rayleigh-Jeans sky temperature (see text): (a) with the old support structure; (b) with the new supports, plotted at the same scale. The smaller error bars in (b) reflect the much greater integration time necessary to detect residual differential ground pickup with the new support structure.

APPENDIX C

Below is a copy of the observing schedule used for the 1993 observations of the SZ effect in Coma. There are four of these schedules: leading and trailing reference fields, as well as shifted versions to equalize the hour angle coverage. This is only the 'Trailing OFF field, Odd scan numbers' version. There are additional versions for the other observing positions in Coma—those schedules are substantially similar.

The suffix `_TO` at the end of the subfiles is required by the scheduler, since all four schedule versions are scheduled at the same time by a master schedule, and since duplicate sub-schedule names are not permitted. The entire preamble to the schedule is included here; this illustrates some of the scheduling cautions that must be heeded.

For details on the scheduler commands and its syntax, see T. Pearson's *OVRO 40m and 5.5m Control System* manual (1993).

```

subfile MEASURE_TO
  flux
  cal
  tag sky
    average
  tag load
    ant load
    ref load
  average
  tag
    ant sky
    ref sky
  zero
endsubfile

subfile ON_TO
  log      starting ON  scan
  obs/wait coma
  sch/sub  MEASURE_TO
  obs/wait t_coma
  record time
endsubfile

subfile OFF_TO
  log      starting OFF scan (trailing)
  obs/wait t_coma
  sch/sub  MEASURE_TO
  obs/wait coma
  record time
endsubfile

```

```

subfile 3C286_TO
  obs/wait 3c286
  flux samples=10 reps=5 idle=10
  cal samples=10
  average
  zero
endsubfile

```

```

! -----
!
! COMA_TO
! -----
!
! Schedule to observe the SZ effect in Coma using a synchronous
! ON-OFF scheme with a TRAILING reference field to subtract ground spillover.
!
! The schedule observes COMA from -3.5 to +4 hrs in hour angle; this
! allows integration with RINGSM observations. This schedule observes 3C286
! at the beginning and at the end of the day. Observing ODD scan numbers.
!
! Each source scan is allotted 15 minutes:
!           4-6 fluxes (depending on the zenith angle)
!           1 cal
!           1 average
!           1 zero_level
!
! A scan on 3C286 takes about 6'20", excluding slew time to the source.
! 3C286 is 15 degrees away from Coma.
!
! The schedule can be started at any time; it will start up correctly, even
! if that means doing an OFF scan.
! In addition, the schedule is designed to recover from scans that take too
! long: if a scan intrudes by more than 30sec into the next (caused by wind
! loading, encoder problems, etc.), the following one will be skipped so that
! the one after it can restart with the correct synchronization. This is
! facilitated by the use of OBSERVE/NOWAITs before each ON/OFF scan subfile.
! The OBSERVE/WAITs within the scan subfiles are retained, along with a
! RECORD TIME to get the LST.
!
! There are 4 versions of this file:
!   Pattern Name   Reference   Scan Config   Schedule
!   -----
!   TO             Trailing   odd           coma_1
!   LE             Leading    even          coma_2
!   TE             Trailing   even          coma_3
!   LO             Leading    odd           coma_4
!
! 04/13/91 first try th
! 04/15/91 improved timing (22s) and added features th
! 04/16/91 added ANWAITs to see when slews are complete th
! 04/21/91 added Cal before starting measurements. th
! 02/11/92 major re-write th
! 03/07/92 changed chart setup th
! 03/09/92 made ON/OFF subfile names unique for T and L th
! 04/18/92 added observe/nowaits before ON/OFF subfile calls th
!           to facilitate synchronization after a failure th
! 04/18/92 improved readability th
! 02/15/93 new season th
!           - A/D cal
!           - all offsets to 0

```

```

! 02/18/93 added scans on 3C286 th
! 02/24/93 refined timing, added records and record time th
! 02/24/93 timing refined again - more idle and more slop th
! 03/30/93 included averages on loads for linearity correction th
!
! -----
!
! setup

hemt on
int 1000 2
threshold 30 3
foff 0 0
off ra=0 dec=0 az=0 za=0
ant mode=sky cal=off heater=off noise=off
ref mode=sky cal=off heater=off noise=off
dicke switched enable
noise 0
atten 6
adcal

default_flux chan=36 azo=-22.16 zao=0 mode=one target=sky
default_cal sam=20 reps=1 idle=3 chan=36 cal=ant target=load
default_ave sam=5 reps=1 idle=3 c0=33 c1=34
default_zero sam=5 reps=1 idle=3

! chart recorder
schedule/file user:[th.chart]chart5m.sch reps=1

! start observing

log -----
log SZ on Coma: ON-OFF, 15min scans, ODD
log TRAILING Reference Field, 2 Scans on 3C286
log START
log -----

record corrections
record offsets

observe/wait 3c286
wait/lst 09:10:00
cal
zero

log --- 3C286 -----
sch/sub/wait/lst 3C286_TO start=09:15:00 stop=09:50:00

log --- Scan 1 ---
default_flux idle=12 reps=6 samples=23
obs coma
sch/sub/wait/lst ON_TO start=09:30:00 stop=09:30:30
obs t_coma
sch/sub/wait/lst OFF_TO start=09:45:00 stop=09:45:30

log --- Scan 3 ---
default_flux idle=12 reps=6 samples=23
obs coma
sch/sub/wait/lst ON_TO start=10:00:00 stop=10:00:30
obs t_coma
sch/sub/wait/lst OFF_TO start=10:15:00 stop=10:15:30

```

```
log --- Scan 5 ---
default_flux idle=13 reps=6 samples=22
obs coma
sch/sub/wait/1st ON_TO start=10:30:00 stop=10:30:30
obs t_coma
sch/sub/wait/1st OFF_TO start=10:45:00 stop=10:45:30
```

```
log --- Scan 7 ---
default_flux idle=15 reps=6 samples=21
obs coma
sch/sub/wait/1st ON_TO start=11:00:00 stop=11:00:30
obs t_coma
sch/sub/wait/1st OFF_TO start=11:15:00 stop=11:15:30
```

```
log --- Scan 9 ---
default_flux idle=17 reps=5 samples=26
obs coma
sch/sub/wait/1st ON_TO start=11:30:00 stop=11:30:30
obs t_coma
sch/sub/wait/1st OFF_TO start=11:45:00 stop=11:45:30
```

```
log --- Scan 11 ---
default_flux idle=21 reps=5 samples=23
obs coma
sch/sub/wait/1st ON_TO start=12:00:00 stop=12:00:30
obs t_coma
sch/sub/wait/1st OFF_TO start=12:15:00 stop=12:15:30
```

```
log --- Scan 13 ---
default_flux idle=28 reps=4 samples=28
obs coma
sch/sub/wait/1st ON_TO start=12:30:00 stop=12:30:30
obs t_coma
sch/sub/wait/1st OFF_TO start=12:45:00 stop=12:45:30
```

```
log --- Scan 15 ---
default_flux idle=30 reps=4 samples=27
obs coma
sch/sub/wait/1st ON_TO start=13:00:00 stop=13:00:30
obs t_coma
sch/sub/wait/1st OFF_TO start=13:15:00 stop=13:15:30
```

```
log --- Scan 17 ---
default_flux idle=24 reps=5 samples=21
obs coma
sch/sub/wait/1st ON_TO start=13:30:00 stop=13:30:30
obs t_coma
sch/sub/wait/1st OFF_TO start=13:45:00 stop=13:45:30
```

```
log --- Scan 19 ---
default_flux idle=19 reps=5 samples=24
obs coma
sch/sub/wait/1st ON_TO start=14:00:00 stop=14:00:30
obs t_coma
sch/sub/wait/1st OFF_TO start=14:15:00 stop=14:15:30
```

```
log --- Scan 21 ---
default_flux idle=16 reps=6 samples=21
obs coma
sch/sub/wait/1st ON_TO start=14:30:00 stop=14:30:30
```

```
obs t_coma
sch/sub/wait/1st OFF_TO start=14:45:00 stop=14:45:30

log --- Scan 23 ---
default_flux idle=14 reps=6 samples=22
obs coma
sch/sub/wait/1st ON_TO start=15:00:00 stop=15:00:30
obs t_coma
sch/sub/wait/1st OFF_TO start=15:15:00 stop=15:15:30

log --- Scan 25 ---
default_flux idle=13 reps=6 samples=22
obs coma
sch/sub/wait/1st ON_TO start=15:30:00 stop=15:30:30
obs t_coma
sch/sub/wait/1st OFF_TO start=15:45:00 stop=15:45:30

log --- Scan 27 ---
default_flux idle=12 reps=6 samples=23
obs coma
sch/sub/wait/1st ON_TO start=16:00:00 stop=16:00:30
obs t_coma
sch/sub/wait/1st OFF_TO start=16:15:00 stop=16:15:30

log --- Scan 29 ---
default_flux idle=12 reps=6 samples=23
obs coma
sch/sub/wait/1st ON_TO start=16:30:00 stop=16:30:30
obs t_coma
sch/sub/wait/1st OFF_TO start=16:45:00 stop=16:45:30

log --- 3C286 ---
sch/sub/wait/1st 3C286_TO

! mark chart recorder and clean up
sch/file [th.chart]mark5m.sch
default_flux idle=10 sam=20 reps=1

log -----
log SZ on Coma: ON-OFF, 15min scans, ODD
log TRAILING Reference Field, 2 Scans on 3C286
log STOP
log -----
```

APPENDIX D

The analysis of single-dish microwave background data is relatively straight-forward. Unlike some observing techniques, such as interferometry, the data entering the analysis are conceptually easily connected to the result—no large-scale transformations have to be accomplished. In addition, the amount of data is small. The total size of all the season's log files for Coma was 2.9MB for 1992 and 3.7MB for 1993, with the increase mainly due to increased daily observing time and additional data recording.

However, the analysis does pose special problems. First, the results are sensitive to subtle software problems: such problems might be extremely hard to recognize and yet affect the result significantly. But more importantly, we demand consistency and repeatability from our analysis. This means that consistent algorithms should be applied with clearly defined parameters. It also means that most, if not all, of the book-keeping should be automatic.

This points to an important feature of our microwave background data: the data consist initially of one single time line of log file records, which must be separated into the parallel strands of diverse information, but which in turn must retain their relationship to each other. For instance, the output of Average procedures to measure the total power level must be kept together with that of Zero procedures that give the instrumental zero level offset. More critically, flux measurements from the ON field must be kept together with those from its corresponding OFF field to ensure proper pairing during ground subtraction. Coupled with the flexibility of possible observing schedules, these considerations require a consistent software approach, while complicating it enormously.

During analysis, it is also vital to be able to 'debug' the data: the ability to produce plots easily and flexibly is critical.

No general software package is available for such a special-purpose application, and none of the standard astronomical analysis packages can be used. This contributed to my decision to write my own analysis software package. Software for the analysis of single-dish SZ data does exist—most notably, the package created by M. Birkinshaw over the years. However, I felt more comfortable creating my own, rather than revising and

adapting another set of programs with which I am not very familiar. This gives the added benefit of cross-checking the two packages for their consistencies.

The development environment chosen for this project has the LISP language as its center; while it was written on a Macintosh computer, it could be converted to other platforms relatively easily, since Common LISP (as defined by Steele 1990) is a standard language and is very portable. The use of LISP has tremendous advantages. Because of its modern nature as a fully object-oriented language, complex programs can be written easily. Its other major advantage is that it is run in an interactive environment (in addition to providing compiled code), in which any code segment can be tested as it is written. This simplifies programming and debugging and helps to fulfill the demand for error-free and robust software.

Even though LISP is a very dense and shorthand language, the analysis package still grew remarkably during its development, mostly in the summer of 1992. As of now, the source code alone has a size of about 500kB.

The analysis package is highly interactive and expandable. It retains the conceptual organization of the data in a hierarchical tree and allows most editing and analysis steps to be performed on any of these elements. As intended, this encourages a careful perusal of the data set.

APPENDIX E

Careful statistical definitions are vital to the result. Even though the equations below are only a restatement of common Gaussian statistics, they are repeated to make the distinctions and the data units clear.

In the analysis, both unweighted and weighted means are used, which are given by the usual

$$\langle x \rangle_{unwt} = \frac{1}{N} \sum_{i=1}^N x_i \quad \text{and} \quad \langle x \rangle_{wt} = \frac{\sum_{i=1}^N \frac{x_i}{\sigma_i^2}}{\sum_{i=1}^N \frac{1}{\sigma_i^2}}.$$

Here, the x_i are the individual data values that are desired in the average. For example, if we are interested in the average value of the CAL over a period of N measurements, they would be made up of the individual Cal procedure outputs. If, however, we are interested in the overall value of our observations of the SZ effect in the center position, for example, the x_i are the values of the N individual data pairs, consisting of the difference of the corresponding individual data points

$$x_i = x_{ON,i} - x_{OFF,i},$$

with the usual error propagation giving

$$\sigma_i = \sqrt{\sigma_{ON,i}^2 + \sigma_{OFF,i}^2}.$$

A σ_i here refers exclusively to the ‘standard error’ $\sigma_{(x)}$ in its associated measurement value, which is related to the standard deviation σ_{x_i} calculated from the measurement’s constituents by the usual

$$\sigma_{(x)} = \frac{\sigma_{x_i}}{\sqrt{N}}.$$

They can be calculated from the constituents of an average as follows in the unweighted case

$$\sigma_{\langle x \rangle}^2 \Big|_{unwt} = \frac{\sum_{i=1}^N (x_i - \langle x \rangle_{unwt})^2}{N(N-1)}$$

and similarly in the weighted one

$$\sigma_{\langle x \rangle}^2 \Big|_{wt} = \frac{1}{(N-1)} \frac{\sum_{i=1}^N \frac{(x_i - \langle x \rangle_{wt})^2}{\sigma_i^2}}{\sum_{i=1}^N \frac{1}{\sigma_i^2}}.$$

The σ_i of a single data point, a single Flux procedure result, is the standard error calculated from the standard deviation of the individual integrations that make up the measurement; as explained above, these integrations have a typical length of 1000ms and are used with equal weight in the procedure's result. The four segments of the Flux procedure are also combined with equal weight by the telescope control software, and the standard error of the result is calculated from error propagation.

We can also calculate the propagated standard error for a weighted mean. It takes on the following simple form

$$\sigma_{\langle x \rangle}^2 \Big|_{wt,prop} = \frac{1}{\sum_{i=1}^N \frac{1}{\sigma_i^2}}.$$

Finally, the moments of a distribution are defined as in Press *et al.* (1986).

APPENDIX F

This appendix contains an outline of the combination of SZ and X-ray data to calculate the distance to the cluster; this discussion closely follows the lines of Birkinshaw *et al.* (1992) and is given here for reference purposes.

The X-ray surface brightness is given by

$$b_x = \frac{1}{4\pi(1+z)^3} \int n_e^2 \Lambda_e dl,$$

while the SZ effect at zero frequency is described by

$$\Delta T = -\frac{2T_e k \sigma_T}{m_e c^2} \int n_e T_e dl.$$

As is plain from the discussion above, the electron temperature and density as well as the energy loss function Λ_e are not constant, but are instead functions of position \mathbf{r} in the cluster. We can express this position in a cylindrical coordinate system with radius θ from a reference line of sight (measured in angular units on the sky), azimuthal angle ϕ about the line of sight, and length along the line of sight ζ , expressed in angular units by $\zeta = l/D_A$, dividing by the angular diameter distance. The positional dependencies can now be rewritten in terms of a reference value for each quantity and a *form factor* f ,

$$n_e(\mathbf{r}) = n_{e,0} f_n(\theta, \phi, \zeta)$$

$$T_e(\mathbf{r}) = T_{e,0} f_T(\theta, \phi, \zeta),$$

with the energy loss function becoming

$$\Lambda_e(\nu, T_e) = \Lambda_{e,0}(\nu) f_\Lambda(\theta, \phi, \zeta),$$

but remaining a function of the observing frequency ν . From these form factors, we can calculate *form integrals*—angles that contain the structural information of the cluster parameters:

$$\Theta_x(\theta, \phi) = \int f_n^2 f_\Lambda d\zeta$$

$$\Theta_{SZ}(\theta, \phi) = \int f_n f_T d\zeta.$$

Inserting those into the above relations and keeping in mind the cosmologically small distance to the Coma cluster, we can easily calculate the Hubble constant from the angular diameter distance:

$$D_A = \frac{cz}{H_0} = \frac{\Delta T_{SZ}^2(\theta, \phi)}{b_x(\theta, \phi)} \frac{\Theta_x(\theta, \phi)}{\Theta_{SZ}^2(\theta, \phi)} \left(\frac{m_e c^2}{kT_{e,0}} \right)^2 \frac{\Lambda_{e,0}}{16\pi T_r^2 \sigma_T^2 (1+z)^3} .$$

If the temperature and density structures and their normalizations are known for the cluster, the Hubble constant can be calculated directly from the observed X-ray and SZ data. This relation is not independent of the electron density, whose value is implicit in the observed X-ray surface brightness, and whose structural dependency is implicit in the form integrals.

Note that the observed parameters as well as the form integrals are functions of the position of the line of sight through the projected cluster; according to the theory those dependencies should cancel each other directly. This provides a powerful consistency check if the data are sufficient to permit it.