

























COMAP Early Science. II. Pathfinder Instrument

James W. Lamb¹ , Kieran A. Cleary² , David P. Woody¹, Morgan Catha¹, Dongwoo T. Chung^{3,4} , Joshua Ott Gundersen⁵,
 Stuart E. Harper⁶ , Andrew I. Harris⁷ , Richard Hobbs¹, Håvard T. Ihle⁸ , Jonathon Kocz^{2,9}, Timothy J. Pearson² ,
 Liju Philip¹⁰ , Travis W. Powell^{1,11}, Lilian Basoalto¹², J. Richard Bond³ , Jowita Borowska⁸, Patrick C. Breyse¹³ ,
 Sarah E. Church¹⁴, Clive Dickinson⁶, Delaney A. Dunne² , Hans Kristian Eriksen⁸ , Marie Kristine Foss⁸ , Todd Gaier¹⁰,
 Junhan Kim² , Charles R. Lawrence¹⁰, Jonas G. S. Lunde⁸, Hamsa Padmanabhan¹⁵ , Maren Rasmussen⁸,
 Anthony C. S. Readhead² , Rodrigo Reeves¹² , Thomas J. Rennie⁶ , Nils-Ole Stutzer⁸ , Marco P. Viero² ,
 Duncan J. Watts⁸ , and Ingunn Kathrine Wehus⁸ 

(COMAP Collaboration)

¹Owens Valley Radio Observatory, California Institute of Technology, Big Pine, CA 93513, USA; lamb@caltech.edu

²California Institute of Technology, 1200 E. California Blvd., Pasadena, CA 91125, USA

³Canadian Institute for Theoretical Astrophysics, University of Toronto, 60 St. George Street, Toronto, ON M5S 3H8, Canada

⁴Dunlap Institute for Astronomy and Astrophysics, University of Toronto, 50 St. George Street, Toronto, ON M5S 3H4, Canada

⁵Department of Physics, University of Miami, 1320 Campo Sano Avenue, Coral Gables, FL 33146, USA

⁶Jodrell Bank Centre for Astrophysics, Alan Turing Building, Department of Physics and Astronomy, School of Natural Sciences, The University of Manchester, Oxford Road, Manchester, M13 9PL, UK

⁷Department of Astronomy, University of Maryland, College Park, MD 20742, USA

⁸Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, N-0315 Oslo, Norway

⁹Department of Astronomy, University of California, Berkeley, CA 94720, USA

¹⁰Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA

¹¹Great Basin Unified Air Pollution Control District, Bishop, CA 93514, USA

¹²CePIA, Departamento de Astronomía, Universidad de Concepción, Chile

¹³Center for Cosmology and Particle Physics, Department of Physics, New York University, 726 Broadway, New York, NY 10003, USA

¹⁴Kavli Institute for Particle Astrophysics and Cosmology & Physics Department, Stanford University, Stanford, CA 94305, USA

¹⁵Departement de Physique Théorique, Université de Genève, 24 Quai Ernest-Ansermet, CH-1211 Genève 4, Switzerland

Received 2021 November 24; accepted 2021 December 18; published 2022 July 13

Abstract

Line intensity mapping (LIM) is a new technique for tracing the global properties of galaxies over cosmic time. Detection of the very faint signals from redshifted carbon monoxide (CO), a tracer of star formation, pushes the limits of what is feasible with a total-power instrument. The CO Mapping Project Pathfinder is a first-generation instrument aiming to prove the concept and develop the technology for future experiments, as well as delivering early science products. With 19 receiver channels in a hexagonal focal plane arrangement on a 10.4 m antenna and an instantaneous 26–34 GHz frequency range with 2 MHz resolution, it is ideally suited to measuring CO ($J = 1-0$) from $z \sim 3$. In this paper we discuss strategies for designing and building the Pathfinder and the challenges that were encountered. The design of the instrument prioritized LIM requirements over those of ancillary science. After a couple of years of operation, the instrument is well understood, and the first year of data is already yielding useful science results. Experience with this Pathfinder will guide the design of the next generations of experiments.

Unified Astronomy Thesaurus concepts: [Astronomical instrumentation \(799\)](#); [Cassegrain telescopes \(202\)](#); [Galaxy clusters \(584\)](#); [Galaxy formation \(595\)](#); [High-redshift galaxy clusters \(2007\)](#); [Molecular gas \(1073\)](#); [Radio astronomy \(1338\)](#); [Radio receivers \(1338\)](#); [Radio telescopes \(1360\)](#)

1. Introduction

The new field of line intensity mapping (LIM; see Kovetz et al. 2019 for a review) promises to open new windows on the universe over a wide range of cosmic epochs. By integrating atomic or molecular line emission from all sources in the line of sight, LIM efficiently affords a view over large volumes of the universe. With this technique, we can probe the evolution of galaxies too faint to be detected using blind pencil-beam surveys. Line emission from the carbon monoxide (CO) molecule has several useful properties for this purpose. With its strong rotational spectrum at millimeter and shorter wavelengths, CO has long been the object of astrophysical

interest for its intrinsic properties and as a tracer of more common but less easily detected molecules, such as molecular hydrogen. Its prevalence makes it a powerful tool for tracking mass associated with star formation (Carilli & Walter 2013; Tacconi et al. 2020). At epochs of interest for early star and galaxy formation redshifted CO emission falls in a spectral window in the centimeter wavelength range that is well suited to ground-based observations. Atmospheric opacity is low, receivers have high sensitivity, radio-frequency interference (RFI) is minimal, and appropriate spatial resolution is achieved with 10 m class antennas.

CO Mapping Project (COMAP) seeks to exploit these advantages to trace the global properties of galaxies through the epoch of galaxy assembly and back to the epoch of reionization using CO LIM (Cleary et al. 2022). As a first step, the Pathfinder instrument targets the 26–34 GHz band, where it is sensitive to CO ($J = 1-0$) (rest frame 115 GHz) from redshift



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Table 1
Principal COMAP Instrument Parameters

| Parameter | Value | Comment |
|--------------------------|----------------------------|--|
| Antenna diameter | 10.4 m | Existing antenna |
| Angular resolution | 4/5 FWHM | $k_{\perp} \lesssim 0.76 \text{ Mpc}^{-1}$ |
| Scan size | $\sim 2^{\circ}$ | $k_{\perp} \gtrsim 0.028 \text{ Mpc}^{-1}$ |
| System temperature | 44 K (median) | Current technology |
| Feeds | 19 | Hexagonal array; |
| Polarization | Left circular | Standing wave control |
| Frequency range | 26–34 GHz | $k_{\parallel} \gtrsim 0.006 \text{ Mpc}^{-1}$ |
| Native freq. resol. | 1.95 MHz | Extended science goals |
| Science freq. resol. | 31.25 MHz | $k_{\parallel} \lesssim 1.5 \text{ Mpc}^{-1}$ |
| Scan rate ^a | 1° s^{-1} | Antenna drive limited |
| Integ. time ^b | 20 ms | Several samples per beamwidth |

Notes.

^a A high scan rate is desirable so that the time to scan the source is small compared to changes in the atmosphere and amplifier gains.

^b This is the fundamental integration time implemented in the spectrometer.

^c Instrument parameters were based on a combination of the science requirements and properties of existing components.

$z = 2.4$ to 3.4 , with a fainter contribution from CO ($J = 2-1$) (rest frame 230 GHz) at $z = 6-8$. COMAP aims to constrain the distribution of CO-emitting galaxies, expressed as a power spectrum that is derived from the Fourier transform of the three-dimensional distribution of sky brightness. The domain of the transformed data is the wavenumber, generally written as k , and the range is the brightness power spectral density.

In brief, the COMAP instrument comprises a 19-feed focal plane array receiver mounted at the secondary focus of a 10.4 m diameter Cassegrain telescope. The 26–34 GHz signals are downconverted to 2–10 GHz for transmission to the antenna sidecab. There they are converted to four 0–2 GHz bands for digitizing, Fourier transformation, and integration in dedicated digital hardware. The telescope is located at the Owens Valley Radio Observatory, which has excellent weather for observing at centimeter wavelengths year-round and full local technical support.

Table 1 lists the main instrumental parameters for COMAP. These are based on available equipment and technology selected to meet the science requirements. Fortuitous availability of a Leighton 10.4 m antenna (Leighton 1977) from the decommissioned CARMA array (Woody et al. 2004) prescribed some constraints, but compromises were minimal, especially compared with the advantage of having an existing highly performant, well-characterized antenna.

Theoretical intensity modeling predictions vary widely and are regularly updated (Chung et al. 2022), but for the purposes of designing the COMAP Pathfinder, a canonical emission spectrum was assumed to be a distribution of sources of about $40 \mu\text{K}$ brightness temperature, with a spectral width for an unresolved galaxy of $\sim 40 \text{ MHz}$ (Li et al. 2016). With a system temperature of $\sim 40 \text{ K}$, the signal-to-noise ratio is of order 10^{-6} . For a spectral channel of width Δf , the standard deviation normalized to the mean for an integration time τ is $\sigma = 1/\sqrt{\Delta f \tau}$. Assuming $\Delta f = 40 \text{ MHz}$, achieving $\sigma = 3 \times 10^{-7}$ requires an integration time of $\tau \approx 3$ days per pixel. A 1 deg^2 patch with $5'$ resolution would then need about 1.3 yr of observing time. Multifield receivers (see, e.g., Emerson & Payne 1995) reduce the integration time in proportion to the number of feeds. A practical configuration of 19 feeds in a

hexagonal close-packed array was adopted for COMAP, reducing the time for this idealized observation to about 24 days of integration time. Multiple feeds also provide extra degrees of freedom to mitigate some systematic errors.

2. Instrument Overview

The selection of 19 feeds was a compromise between integration speed and cost. Moreover, the speedup is less than proportional to the number of feeds since aberrations and blockage start to become significant limitations.

Coverage in k_{\perp} -space is determined by angular resolution and scan extent, while the k_{\parallel} -space range is defined by the frequency resolution and span. Table 1 shows the k ranges defined for the COMAP instrument. These fit well with the target k -range of at least $0.05-0.5 \text{ Mpc}^{-1}$, a reasonable choice based on current modeling (Chung et al. 2022). The scan extent is not limited by the antenna, but by covering a small field long enough to achieve a statistically significant measurement based on the fiducial model. In the spectral domain the resolution does not need to be much narrower than the galaxy or galaxy cluster line widths, but other science, such as radio recombination line detection, is enabled by the 1.95 MHz resolution.

The following sections provide an overview of the design and implementation of the COMAP instrument. In the design details, we focus on aspects that are particularly challenging for LIM. These are viewed in the context of general considerations of the stability of subsystems over timescales of interest. We present measurements of components, subsystems, and finally the overall performance of the system based on observations of blank sky and astronomical objects such as Jupiter. In the discussion section we use the experience gained to extrapolate to future instruments.

In designing the COMAP system it was possible to focus on the challenges posed by LIM without regard for requirements for more general observations. Stability of the passband was one of the most critical issues, and any potential perturbing effects, particularly standing waves, were evaluated and mitigated as much as possible. Adverse influences are primarily temperature changes and cable flexing. In evaluating the importance of standing waves, we need to know the ripple amplitude in the spectrum relative to the nominal power, P_r , the mean frequency in the spectrum, f_{mean} , the integration time per pixel, t_{int} , the scan time, t_{scan} , and the resolution bandwidth, Δf . In Appendix A.1, we derive a criterion that ensures that the rms of the *change* in ripple over the scan time is less than the statistical noise if

$$\frac{d\tau}{dt} \ll \frac{2}{\pi f_{\text{mean}} P_r t_{\text{scan}} t_{\text{int}}^{1/2} \Delta f^{1/2}}, \quad (1)$$

where τ is the change in the two-way transit time in the standing wave “cavity.” For cases where this is thermally driven, this can be used to bound the acceptable temperature change rate, dT/dt .

Spillover and sidelobes falling on the ground or bright celestial objects are other potentially detrimental effects. Strategies to avoid these phenomena are discussed in the relevant component or system sections. Details of how these effects were evaluated are given in the Appendix.

We will outline the general physical layout of COMAP and the signal flow from the antenna to the spectrometers, before presenting more detailed descriptions of constituent components.




Figure 1. Main components of the COMAP telescope. A new secondary mirror was fabricated so that the focal plane receiver could be placed in front of the primary (see text for details). Another major modification was the “pop-out” extension to the sidecab for the digital electronics. The two air conditioning units critical to performance are indicated.

2.1. Telescope Layout

Figure 1 shows the Leighton 10.4 m antenna with the receiver package mounted at the secondary focus. Matching the primary to the focal plane array receiver required a new secondary mirror that was larger than the original, as well as having a shorter focal length. An ambient temperature absorber vane is mounted on the receiver package so that the system noise temperature can be periodically calibrated. The receiver package has a cryostat with a mechanical cryocooler for the low-noise amplifiers (LNAs), and its compressor is located on the azimuth platform, connected with flexible helium lines. Signals from the receiver are sent on coaxial cable to the sidecab where the back-end electronics are located.

To accommodate the required volume of electronics in the sidecab, a “pop-out” was added. Most of the pop-out is occupied by an RFI-tight cabinet that contains the digital electronics with its own air conditioning. The top of the pop-out and the remainder of the sidecab are used for analog electronics, antenna control computer, and the elevation cable wrap. It uses the original sidecab air conditioning and temperature control. Connections to the central electronics building for frequency references, control and monitoring, and data transfer are on optical fiber.

2.2. Signal Path

Figure 2 is a schematic view of the radio-frequency (RF) path, starting from the feed horns. These are attached to the outside of the cryostat with waveguide vacuum window assemblies. Cooled polarizers convert left circular polarization (LCP) to the linear polarization accepted by the waveguide 26–34 GHz LNAs. Signals from the LNAs are mixed in the first downconverter modules (DCM1) with a 24 GHz local

oscillator (LO), converting them to an intermediate frequency (IF) of 2–10 GHz. There are four DCM1s, each processing five receiver outputs. In addition to the frequency conversion, they amplify the IFs before transmitting them on coaxial cables to the sidecab.

In the sidecab, the signal from each feed is again amplified and split into band A (2–6 GHz) and band B (6–10 GHz). Each band goes to a second downconverter module (DCM2) with an in-phase/quadrature (IQ) mixer fed by an LO for the relevant band (4 and 8 GHz, respectively). In a conventional sideband-separating system (Maas 1992), the in-phase (I) and quadrature (Q) outputs at 0–2 GHz would be connected to an analog quadrature hybrid that recombines the signals into the lower sideband (LSB) and upper sideband (USB) of the relevant LO. Cleanly separating the sidebands demands extremely good phase and amplitude balance, which is difficult to achieve over a wide bandwidth. A digital quadrature combination is significantly more accurate (Finger et al. 2013). The I and Q signals are directly digitized in ROACH-2 spectrometers (CASPER Collaboration; Parsons et al. 2008; Digicom Electronics, Inc.) and converted to complex spectra with a polyphase filter bank (PFB; Price 2021). I and Q spectra are combined with complex weights on a channel-by-channel basis to give well-separated USB and LSB spectra, and the squared magnitudes are integrated for 20 ms. Each receiver channel then has four output data streams corresponding to sky frequencies denoted by A:LSB (26–28 GHz), A:USB (28–30 GHz), B:LSB (30–32 GHz), and B:USB (32–34 GHz). These data streams are captured on a high-speed Ethernet switch and forwarded to the data storage computer.

3. Instrument Design Details

3.1. Antenna

One of the Leighton 10.4 m antennas (Leighton 1977) became available at an opportune time for COMAP. These antennas have been used successfully up to 270 GHz, guaranteeing excellent performance at 30 GHz. Significant modifications were required for COMAP to accommodate a focal plane centimeter-wave receiver, including an optical redesign requiring a different secondary mirror and a new location and mount for the receiver. The digital electronics were located in the antenna sidecab, necessitating a modest expansion of its volume. In previous use at the CARMA observatory, the antenna control system was part of a highly integrated software system that was not readily reusable for COMAP. However, it was possible to isolate the antenna drive code running in the antenna and bind it to another control system developed at OVRO.

3.2. Optical Design

Figure 3 shows the optical parameters for the COMAP telescope, and the values are listed in Table 2. Previous receivers were located behind the primary mirror or in the sidecab coupled with a beam waveguide. For a wavelength of 1 cm the clearance in the backing structure is insufficient for multiple beams to pass through, so the receiver package had to be mounted in front of the surface of the primary, requiring a new secondary mirror with a shorter focal length. This also reduced the required aperture size for the feed horns. To reduce the horn size further, the secondary diameter was increased from 600 to 1100 mm. That diameter is within the size of the aperture at the center of the primary and so incurs no additional blockage penalty.

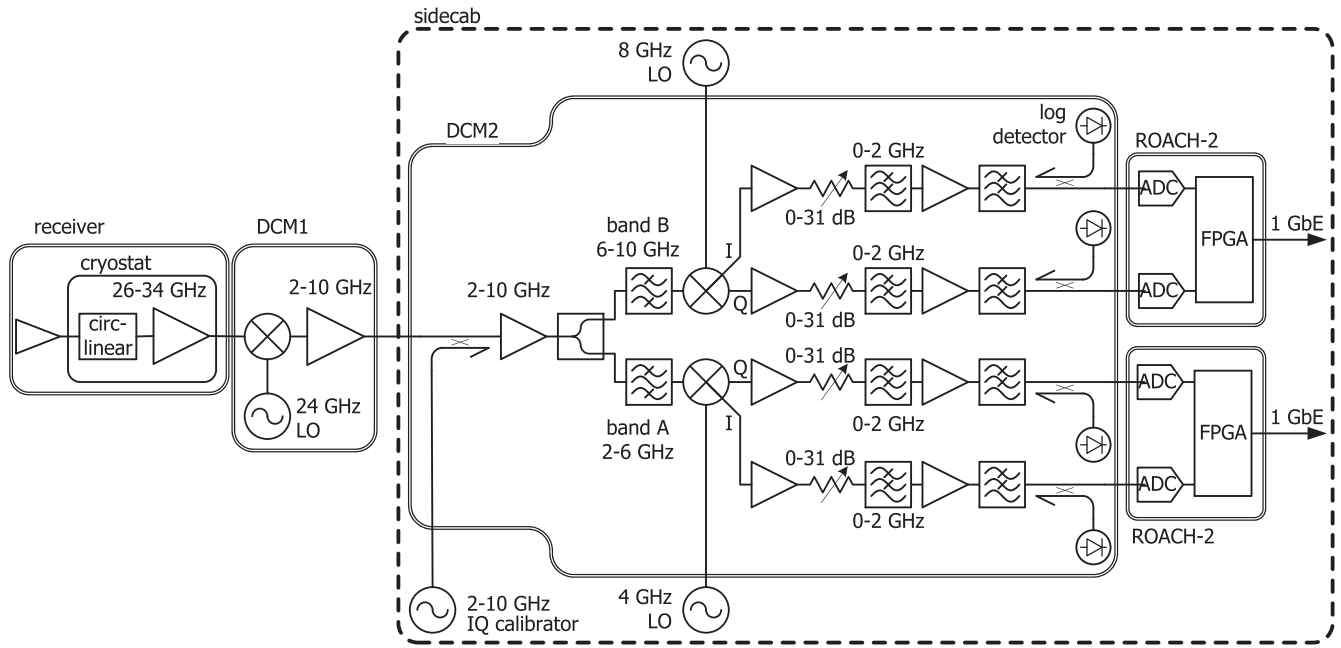


Figure 2. Signal path for the COMAP receivers. This path is essentially identical for each of the 19 pixels, except that there is one 24 GHz LO per DCM1, and one 4 GHz LO, one 8 GHz LO, and one IQ calibration source common to DCM2s for all pixels.




Figure 3. Optical design parameters for the COMAP antenna. Values of the parameters are listed in Table 2.

Table 2
Antenna Optical Parameters

| Parameter | Symbol | Value | Units |
|---|------------|----------|-------|
| Primary mirror diameter | d_p | 10.400 | m |
| Primary focal length | f_p | 4.123258 | m |
| Primary focal ratio | f_p/d_p | 0.396 | |
| Secondary to receiver distance | d_{sr} | 2.114 | m |
| Half angle subtended by primary | θ_p | 64.469 | deg |
| Magnification | M | 4.50 | |
| Effective focal length | $f_s = Mf$ | 18.555 | m |
| Secondary eccentricity | e | 1.571429 | |
| Full secondary mirror diameter ^a | d'_s | 1.300 | m |
| Actual secondary mirror diameter ^a | d_s | 1.100 | m |
| Half angle subtended by secondary | θ_s | 13.847 | deg |
| Secondary interfocal distance | f_c | 2.58426 | m |
| Hyperbola parameters | a | 0.82227 | m |
| | b | 1.20823 | m |
| | c | 1.29213 | m |

Notes. Values for the optical design. See also Figure 3.

^a Full secondary mirror defines geometry. Actual secondary is undersize to reduce spillover round primary.

Choosing a short secondary focal length keeps the feeds and receiver package small, but if it is too short, aberrations are increased, and the blockage of the spherical field between the primary and secondary becomes significant. Additionally, the receiver support becomes unmanageably tall. With the parameters in Table 2, the receiver is as close to the secondary as possible without affecting the blockage.

Increasing the diameter of the secondary also reduces spillover from edge diffraction, which varies as the inverse square root of the diameter in wavelengths. By undersizing the secondary diameter by 200 mm, some of the diffracted wave falls on the primary and is reflected to the sky, further reducing ground spillover. Some slight difference in the spillover among feeds will result from differences in the feed offset from the optical axis.

Manufacturing of the secondary mirror was done in an external machine shop on a numerically controlled mill from a billet of 6061 alloy aluminum. The rear surface had pockets milled to reduce the weight without compromising rigidity, and the final mass was 29.2 kg. The hyperbolic surface has a thickness of 3 mm.

As we discuss in Appendix A.1, standing waves that cause ripple in the passband are a significant problem, so a diverter cone was added to the center of the secondary mirror to reduce the optical standing wave. The voltage standing wave ratio (VSWR) was optimized with a wave analysis using a closed-form solution (Padman & Hills 1991) and the resulting design verified using TICRA Tools physical optics (PO) analysis.¹⁶ The cone, featured

¹⁶ <https://www.ticra.com/software/grasp/>

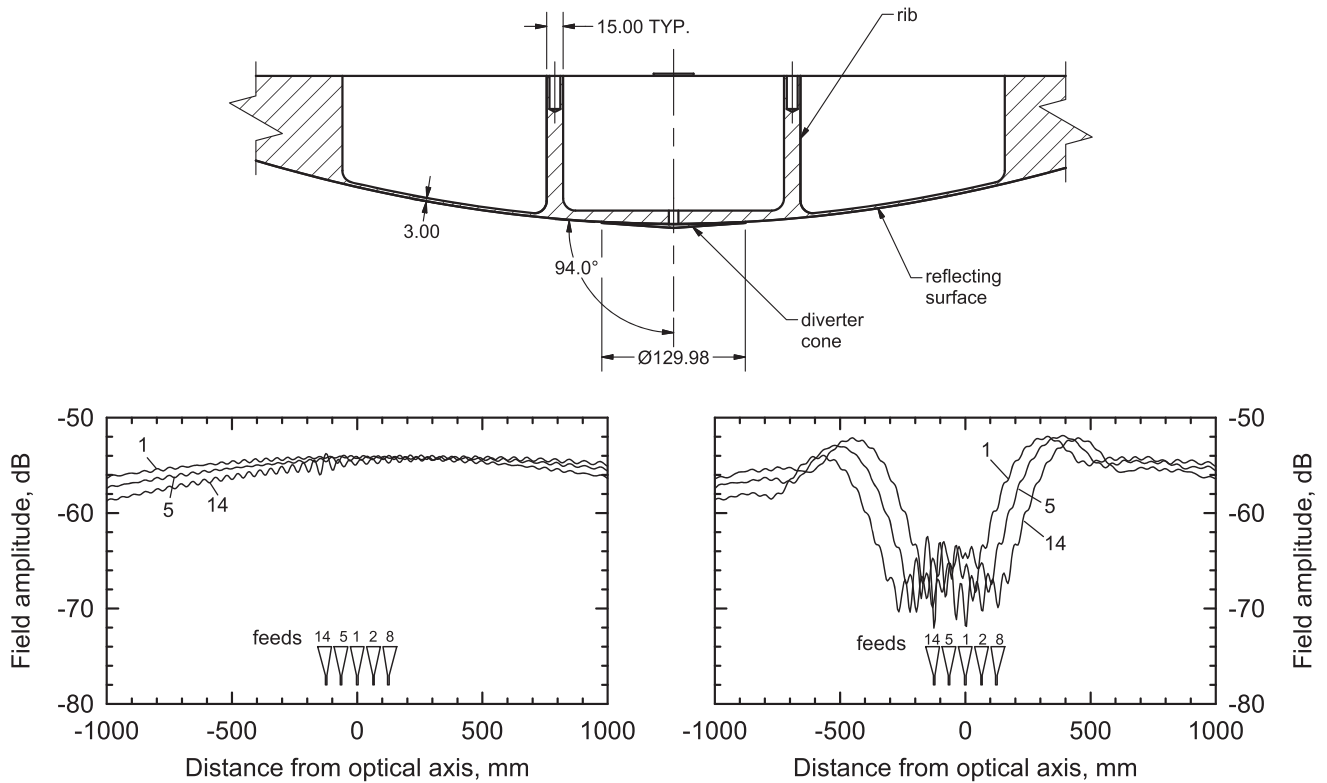


Figure 4. Top: cross-section of the center of the secondary mirror, showing the diverter cone designed to reduce the optical VSWR. Dimensions are in millimeters. Bottom: calculated field amplitude at 30 GHz in the feed aperture plane after reflection by the secondary mirror with the diverter cone (left: no diverter cone; right: with diverter cone). The fields from representative feeds (1, 5, and 14; see Figure 5) are plotted. The feeds drawn to scale at the bottom of the plot illustrate what fraction of the reflected field is captured from the same or a different feed, indicating self- or cross-coupling.

in Figure 4, is very shallow and approximately tangential to the hyperboloid where they meet. This cone was machined separately from the mirror and has a threaded stem that screws into the center of the secondary.

By calculating the field from the horn when it is scattered back to the horn aperture plane, the effectiveness can be evaluated. Figure 4 shows that the diverter cone reduces reflections by about 10 dB.

Corrugated horns (Clarricoats & Olver 1984) were selected as feeds since they have excellent match and beam patterns up to at least 40% bandwidth. The frequency dependence of the beam pattern is controlled by the phase variation across the aperture, which depends on the length and flare angle of the horn (Thomas 1978). A small phase deviation gives a diffraction-limited beam with a width proportional to wavelength, and a large phase deviation produces a constant-width beam. While a constant-width beam maintains aperture efficiency over the band, a diffraction-limited feed gives a frequency-independent antenna beam. The latter was deemed more appropriate for COMAP since it gives a similar resolution in k_{\parallel} at all frequencies. Achieving the desired diffraction-limited feed pattern would require an excessively long horn, but a reasonable compromise was found that gave a modest variation of beamwidth with a 220 mm long feed.

Ideally, the feed apertures would be located on the spherical Petzval surface (Born & Wolf 1980) and pointed near the center of the secondary. However, to simplify the mechanical layout, all the feed apertures were located on a common plane. This has a minimal impact on performance since it has only a small effect on the main beam. The 19 feed apertures are

arranged in a hexagonal close-packed array with a spacing of 65 mm between axes and 1 mm between the edges (Figure 5).

Detailed performance calculations were carried out with TICRA Tools to quantify spillover, aperture efficiency, beam spacing, and sidelobes. PO plus physical theory of diffraction (PTD) calculations (Rudge et al. 1986) were applied to a detailed model of the telescope, including the panels, secondary with diverter cone, and secondary supports. Beam patterns calculated for the horns using CWGSCAT (Hoppe 1993) were used to define the source fields.

The horn dimensions selected in combination with the secondary design result in a telescope FWHM beamwidth of about $4'5$ and a spacing on the sky of $12'0$. As shown in Figure 6, the beamwidth varies by about 4% over the observing band, significantly less than would be expected for a diffraction-limited aperture antenna.

Differences in aperture efficiency for the various feed locations are shown in Figure 7. These are small and due mainly to coma. Achieving maximum efficiency requires a frequency-independent illumination of the telescope aperture, so the decrease with frequency is expected and not detrimental for this experiment.

Sidelobes were characterized on a number of scales and at various frequencies using the TICRA Tools package, with a typical map shown in Figure 8. Element sizes for the finite-element analysis were of order a half wavelength so that the computed field is valid over the forward hemisphere. Limitations are due mainly to some necessary simplifications of the structural elements. For example, the secondary support struts were modeled as two-dimensional elements corresponding to the faces closest to the axis of the dish.




Figure 5. Layout of the feed horns as viewed from the front of the receiver. The numbers at the centers of the feeds are traced through to the final maps, while the indices at the top and side are used for the antenna analysis software. Dimensions for the feed separation and feed horn outer diameter are indicated in millimeters. In the current installation on the telescope, the receiver was rotated by 90° to this orientation.




Figure 6. Calculated beamwidth as a function of frequency for all the beams on the sky. The $1/f$ dependence represents the behavior of a diffraction-limited beam.

Several sources of sidelobes are identifiable. Diffraction by the antenna aperture produces near-in sidelobes that rapidly fall off from the main beam. Spillover of the feed pattern past the secondary results in an annular sidelobe at about 14° from boresight. From the work of Rusch et al. (1982), the scattering of the outgoing plane wave from the primary by the secondary support struts is predicted to propagate along a conical profile, with the axis of the cone along the strut, and the cone containing the optical axis. This produces annular sidelobes that intersect at the main beam and have a width corresponding to the diffraction size for the length of the strut projected on the aperture plane. Wide-angle rear lobes result from the feed pattern being diffracted at the rim of the secondary and then propagating back past the edge of the primary. Central blockage by the hole in the primary creates a (negative) diffraction pattern similar to the primary, but with a scale size 10 times greater and an amplitude of a few percent.

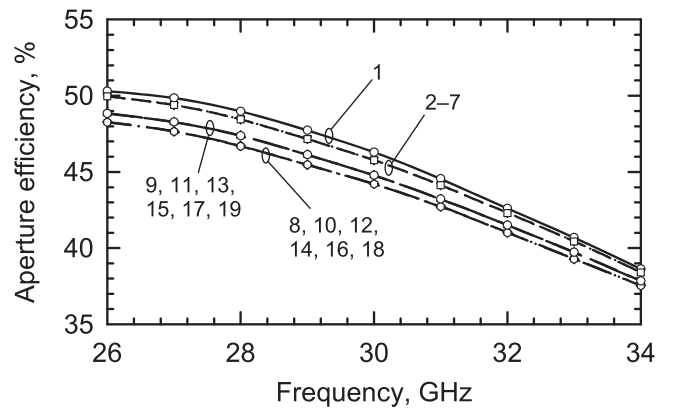


Figure 7. Predicted aperture efficiency, referred to the feed aperture plane. The numbers denote the feeds represented by each set of curves. Although there are only four different offsets from the axis, there are other effects, such as the fourfold feed support symmetry, that make nominally equivalent feeds different. These differences are barely discernible on the graph.

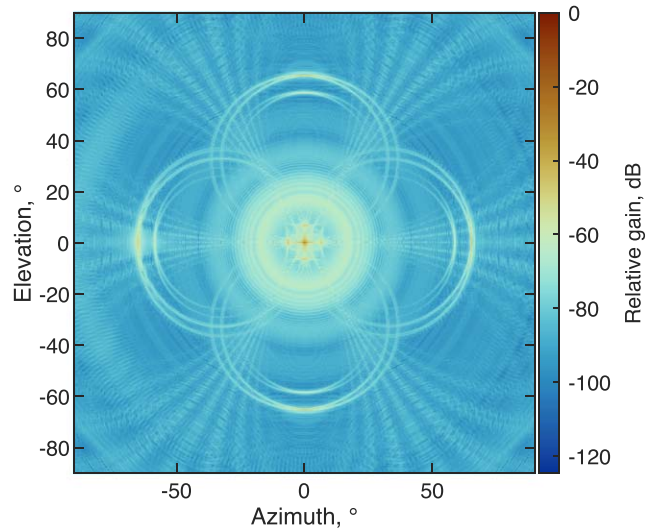


Figure 8. Map of far-out sidelobes. The main beam is about 1 pixel. Sidelobes are due to diffraction at the aperture, spillover past the secondary, secondary support struts, and panel edges.

3.3. Receiver Mount

To mount the receiver in front of the primary, a hexapod made from aluminum rods was clamped at the lower end to struts in the primary backing structure, extending through the central hole in the primary to the aluminum support plate for the cryostat (Figure 1). The triangular configurations match the sixfold symmetries the backup structure is based on and form a strong support without overconstraint. The orientation was chosen to make the cryocooler cylinder parallel to the elevation axis to minimize potential gravitational effects on temperature stability.

Aluminum struts were selected so that the differential thermal expansion relative to the steel antenna structure would keep the feed-to-secondary distance approximately independent of ambient temperature, minimizing variation of the optical standing wave ripple. For a homogeneous steel structure, the temperature sensitivity of the feed-to-secondary distance would be $25 \mu\text{K}^{-1}$. If we apply Equation (1) with $f_{\text{mean}} = 30 \text{ GHz}$, $P_r = 0.01$, $t_{\text{scan}} = 20 \text{ s}$, $t_{\text{int}} = 100 \text{ ms}$, and $\Delta f = 32 \text{ MHz}$, and we




Figure 9. Sectioned view of receiver. The outer vacuum jacket is machined from two aluminum blocks, with aluminum end plates. Details of the waveguide components and the thermal isolation are provided in the text.

require the ripple to be an order of magnitude less than the statistical noise, the rate of change of temperature should be less than 36 mK s^{-1} . Changes in ambient temperature of this order are expected, though the thermal mass of the struts reduces this somewhat. By combining aluminum struts with the steel, a significant improvement is predicted. Since the coefficient of thermal expansion of aluminum is about twice that of steel and the aluminum components span about 2 m, while the distance in steel is about 4 m, the net expansion of the feed-to-secondary distance is zero to first order.

A hexagonal shroud fabricated from aluminum sheet, painted white, and lined with 25 mm thick building foam provides substantial isolation from convective and radiative heat exchange due to ambient temperature and solar radiation. Air is blown into the shroud through a flexible coaxial aluminum duct from the temperature-controlled sidecab. The diameters of the inner and outer ducts are 150 and 200 mm, respectively, and the space between is filled with insulating foam beads. A second identical duct returns the air to the sidecab. The supply duct also carries the signal cables from the receiver down to the back-end electronics.

3.4. Drives

No modifications were made to the existing antenna drive motors or controllers since they had adequate performance. DC motors with brushes are driven by commercial electronic controllers that servo the rate to an applied voltage using

tachometers integrated with the motors. Elevation backlash is removed by unbalancing the elevation structure slightly so that the ball-screw drive is always in compression. Azimuth backlash is countered by having two motors with opposite but unequal torque driving a large azimuth gear.

3.5. Receiver Package

Nineteen feed horns are mounted on the front of a cryostat (Figure 9) that contains circular-to-linear polarizers and cryogenic LNAs. DCM1 modules attached to the outside of the cryostat convert the signal from 26–34 GHz down to 2–10 GHz for transmission to the second stage of signal processing in the sidecab. A microwave-transparent window mounted on the top opening of the shroud (Section 3.3) protects the receiver.

3.5.1. Cryostat

Machined from solid aluminum, the main vacuum container of the cryostat is 375 mm high with end plates installed and 381 mm between flats of its hexagonal cross section. The top face has mounting holes for the vacuum windows that mate with the feed horns. Inside the vacuum space, an aluminum heat shield is attached to the top plate with G-10 fiberglass-epoxy tabs. These are close to the perimeter of the shield, with their face normals directed to the cryostat axis so they can flex on cooling while keeping the shield coaxial with the outer

container. A gold-plated second-stage copper shield is mounted in a similar way inside the first stage shield. The receiver components are individually supported on their respective vacuum window flanges, and thermal connections to the second stage are made with C10100 copper braid, allowing excellent heat transfer with minimal stress (particularly from differential contraction on cooling). Signals from the LNAs are routed using 0.085 stainless steel coaxial cables to a bulkhead on the first stage, and 0.141 stainless steel coaxial cables from there to the vacuum wall vacuum feedthroughs.

Cooling is provided by a Sumitomo RDK-408S2 cold head operating with a CSA-71A compressor. The cold head is inserted through one side of the outer wall of the cryostat, and the first and second stages are attached with multiple sheets of high-purity copper foil to their respective radiation shields.

3.5.2. Calibration Load

Following standard millimeter-wave methods (Penzias & Burrus 1973; Ulich & Haas 1976), amplitude calibration uses an ambient absorber that can be rotated into place in front of the feeds in about a second. The TK RAM absorber (Thomas Keating Ltd) used has a reflection of better than -25 dB at 30 GHz. Tessellated 100 mm square tiles were trimmed to a circular perimeter and mounted on a lightweight frame, and a synchronous AC gear motor inserts and removes it from the beams. With an appropriate geometry for the motor axis, the 432 mm diameter vane rotates from parallel to the window in the calibration position, to being in the plane of the secondary supports in the “out” position (see Figure 1). In that orientation it causes no additional shadowing of the aperture.

Thermal insulation foam is glued to the rear to minimize thermal gradients, and both front and back are sprayed with white paint to reduce solar heating. Temperature sensors are embedded in the load, and the control system continuously records the value. During early testing, it was found that performance was improved by backing the load with aluminum foil. This eliminated the radiation transmitted from the sky and effectively doubled the optical depth of the absorber.

3.5.3. Waveguide Components

Figure 10 shows the chain of waveguide components in the receiver. The corrugated horn was optimized for beam pattern and return loss using CWGSCAT (Hoppe 1993). Pattern measurements in an anechoic chamber were in excellent accord with the predictions as is evident in Figure 11.

Since corrugated structures have very low losses (Clarricoats & Olver 1984), the horn is not cooled. Any slight increase in noise from ambient temperature loss in the horn is offset by avoiding the loss in thicker windows required for cooled horns. This configuration also allows a very small thermal radiation area between ambient and cryogenic temperatures, defined by the area of the waveguide. Vacuum integrity is ensured by a $25\ \mu\text{m}$ thick DuPont Mylar (PET) window, located on the inside of the vacuum flange, allowing feed horns to be removed without breaking vacuum.

As part of the effort to control ripple in the system, the receivers were designed to use LCP. Noise emitted by the LNA and reflected by the secondary has its polarization reversed so it does not couple back to the LNA and create ripple in the noise temperature. Similarly, LCP signals from the sky that are reflected by the LNA and secondary are rejected. Right circular




Figure 10. Exploded view of the waveguide components that are located inside the cryostat. The fiberglass sleeve is attached to the polarizer at the left end only, supporting it so that the right end is about 0.5 mm from the choke groove on the window flange (see inset, top left). The window flange is at ambient temperature and captures the vacuum window against the O-ring seal.




Figure 11. Computed and measured feed patterns. The patterns were measured in an anechoic chamber.




Figure 12. Cutaway view of circular-to-rectangular polarizer. The vanes that provide the retardation are silver-epoxied into slots in the circular waveguide.

polarization from the sky that is reflected at the rectangular waveguide and then by the secondary as LCP couples into the receiver but is uncorrelated with the direct LCP signal and does not produce ripple.

Conversion from circular to linear polarization is effected in a two-stage polarizer, illustrated in Figure 12. The principles of this design are given by Plambeck & Engargiola (2010) based on work by Pancharatnam (1955). It uses waveguide sections that retard one linear polarization relative to the other. The orientation of the principal axes and the retardation per section are chosen to optimize performance across a given band. Increasing the number of stages improves performance but requires increasingly tight tolerances. Electromagnetic modeling indicated that good performance can be achieved with two stages.

Figure 12 shows the construction of the polarizers. In each section, retardation of one component of polarization results from two copper fins symmetrically inserted into slots in a brass circular waveguide. The fins are tapered to reduce mismatches and silver-epoxied (EPOTEK H20E) to the waveguide with a jig to ensure alignment. Different lengths and orientations of the fins are used for the two stages. To minimize the overall length, a small separation is used between the stages that requires careful optimization in the simulator to account for interstage effects. For corrosion protection and loss reduction a soft-gold plating layer is applied. Verification of the manufactured components was done with two polarizers using a vector network analyzer. This verified that the reflection coefficients were better than -20 dB across the band, and the through signal was found to be essentially independent of the relative rotation of polarizers about their common axis.

As seen in Figure 10, a G-10 fiberglass-epoxy cylinder holds the polarizer at the cryogenic end, while the other end is anchored to a flange at ambient temperature. A choke ring (Helszajn 1978) on the ambient flange prevents leakage through the gap between it and the polarizer. On cooling, the gap becomes smaller owing to the different thermal coefficients of expansion between the brass polarizer and the fiberglass epoxy. Simulations confirm that the choke ring operates correctly over the range of gap sizes from ambient to cold. Apertures were machined into the fiberglass-epoxy cylinder to reduce thermal loading. That reduces the stiffness, potentially

allowing misalignment between the polarizer and the choked flange, so the deformation under gravity was calculated using finite-element simulation and the resultant deformed structure was simulated electromagnetically. The effects were determined to be negligible. Measurements of the deflection and transmission of a fabricated device closely matched the simulations.

An electroformed transition was designed to mate the circular polarizer waveguide to the LNA. The mandrel was an aluminum cylinder with flat faces tapered down on four sides to the required rectangular waveguide cross section at the end. Machining and electroforming were done commercially (A. J. Tuck Co.). Anticocking flanges (Kerr et al. 1999) were used for good integrity of interfaces.

Full electromagnetic simulations of the structure from the LNA input to the horn aperture were carried out using Ansys HFSS¹⁷ to verify performance. This included the vacuum window and the waveguide gap choke. Leakage through the gap was verified to be insignificant in terms of cross-talk between feeds.

3.5.4. Low-noise Amplifiers

The front-end amplifiers set the noise performance of the whole receiver provided that they have sufficient gain. Amplifiers based on high electron mobility transistors (HEMTs) provide the lowest noise at Pathfinder frequencies. For array applications, LNAs constructed from monolithic microwave integrated circuits (MMIC) can provide consistent performance without the need for extensive fine-tuning. We therefore selected a proven MMIC HEMT-based LNA design for COMAP (Tang et al. 2006) as the basis for our Pathfinder LNA module. This design, fabricated on Northrop Grumman Corporation’s (NGC) 100 nm InP process, was demonstrated to exhibit noise temperatures in the 7–15 K range over 26–34 GHz. A pool of such devices was available at JPL, and we also fabricated a new version of this design on NGC’s 35 nm InP node. Both 100 nm and 35 nm chips were used interchangeably in the final LNA modules.

The module input is rectangular WR-28 waveguide in order to mate directly with the preceding rectangular-to-circular waveguide transition, as described above. Inside the module (see Figure 13), a waveguide probe couples the incoming RF signal to microstrip, and from there it enters a chain of two MMIC amplifier chips isolated from each other with a 3 dB attenuation pad. The amplifiers are followed by a high-pass filter designed to cut off the first LO frequency of 24 GHz and prevent feedback of any such LO signal that leaks back into the front end. The output of the filter is then coupled to a coaxial bulkhead feedthrough. Bias for the amplifiers is conditioned through a printed circuit board inside the module. Thermally isolating cables conduct the signals to the cryostat wall (Section 3.5.1).

Each LNA module was characterized at cryogenic temperatures; Figure 14 shows the resulting gain and noise temperatures for the 19 modules used in the Pathfinder. The median values over the band were in the range of 39.9–46.4 dB and 8.1–14.4 K, respectively, meeting the system requirements.

¹⁷ <https://www.ansys.com/products/electronics/ansys-hfss>




Figure 13. Left: interior of the COMAP LNA module showing the chain of RF components and the bias conditioning PCB. Right: module exterior showing the waveguide input.




Figure 14. Gain (top) and noise temperatures (bottom) of a representative group of 19 LNA modules.

3.5.5. First Downconverter (DCM1)

Four shielded die-cast enclosures are mounted on the sides of the cryostat, each housing five downconverters. A 24 GHz dielectric resonator oscillator (DRO; Luff Research, PLDRO-24000-10) in each enclosure is phase locked to the 10 MHz source and drives five Marki ML1-1644IS mixers, and RF Lambda RLNA02G10GA amplifiers boost the signal for transmission via 15 m cables to the sidecab. Heatsinks with fans on the outside of the modules conduct heat from the amplifiers. Each module has a control interface board that allows the amplifiers to be turned on or off and to monitor the temperatures of the amplifiers.

3.6. Sidecab

Originally the sidecab accommodated millimeter receivers and the antenna control computer, and the signals were sent back to a central building for processing. For COMAP it was unclear what effects would be caused by a long optical fiber transmission to a remote building. Initial tests revealed some instability, but that was likely due to cycling of the air conditioning in the building. In any event, the decision was

made to put the digital electronics in the antenna sidecab. The millimeter receivers were removed to allow for an elevation cable wrap, but the antenna control rack was left in place. The original thermal control, comprising an air conditioner and a heater controlled by a commercial regulator and a blower, keeps the sidecab temperature steady at 21°C.

Due to the volume of electronics, a “pop-out” extension was added (Figure 1). An RFI-tight cabinet in the lower section contains the digital systems, and a split system air conditioner was installed in it to handle the additional heat load. The compressor located on the antenna platform has a variable-speed drive that maintains a very constant temperature. A small modification was required to override one of its modes that causes significant temperature cycling during winter months.

DCM2 modules are housed in a card cage on top of the rack, with a fan to blow the temperature-regulated sidecab air over them. LMR-240 cables connect the DCM2 outputs to a panel on top of the digital, and then to the ROACH-2 spectrometers.

3.7. Cable Management

Fifteen-meter-long LMR-400 cables run from the DCM1 modules on the cryostat exterior to the DCM2 modules in the sidecab. This choice was based on having a cable with acceptable loss variation across the 2–12 GHz IF band and a compromise between stiffness and flexibility to accommodate the elevation cable wrap without being too floppy. As mentioned above, the cables are routed from the receiver to the sidecab via an insulated duct. Once inside the sidecab, they are individually clamped to one straight edge of a plate curved with a 380 mm radius. Passing over the plate, they hang off the other end in loops that attach to the sidecab wall. As the antenna changes elevation, the curved plate rotates with the tipping structure and the hanging loops accommodate the change in elevation angle.

From the clamps on the sidecab, the cables are carried by a cable tray over to the DCM2 modules at the other end of the sidecab.

Although ripple from this length of cable corresponds to k values outside the range of interest, combining the 2 MHz samples into 31.25 bins aliases this ripple from a period of 8.5 MHz ($k_{\parallel} \approx 5.5 \text{ Mpc}^{-1}$) to 100 MHz ($k_{\parallel} \approx 0.5 \text{ Mpc}^{-1}$). While this is in the relevant range, binning reduces the amplitude. Additionally, the cable loss and the 3 dB matching pads at either end keep the ripple well below 1% and, with the flex and temperature controls, make it a negligible effect.

3.8. Second Downconverter (DCM2)

For each receiver, two downconverter boards are mounted on an aluminum plate for the two bands, A and B. The cables from the DCM1s go to SMA connectors on the front of the plate, generally with a coaxial attenuator selected to ensure that the following IF amplifier (RF Lambda RLNA02G10GA) is in its linear range. After the amplifier, a 16 dB coupler allows a calibration tone to be injected, and then the signal is split in two, going to band-defining filters for bands A and B (Figure 2). Identical downconverter boards are used for the two bands.

Each board has a Marki MLIQ-0218SM IQ mixer followed by gain stages and digital attenuators (0.5 dB steps, 31 dB range) and log detectors for both I and Q. Two 0–2 GHz low-pass filters are included to reduce aliasing and attenuate LO

feedthrough. At the output a 3 dB pad is used to improve the match to the ADC circuit in the digital processor that is connected with LMR-240 cable.

LOs are either a 4 GHz DRO (Luff Research PLOA4000-10) or an 8 GHz DRO (Luff Research PLOA8000-10). Each is phase locked to the station 10 MHz reference clock, and their outputs are distributed in a series of stripline power splitters.

The critical component for linearity in the system is the IF amplifier on the DCM2 plate. For each receiver, a pad between the cable and the DCM2 plate was chosen to ensure that the operating levels were in the linear range at the 1% level when the ambient load was in place. Because of the variation in power from the receivers, there is a spread of about 10 dB in pad values.

Cables from the DCM2 to the digital rack, and then from the rack to the ROACH-2s, have total lengths up to 2.6 m. For a typical cable temperature sensitivity of $5 \times 10^{-6} \text{ K}^{-1}$ Equation (1) requires stability on the order of 3 K s^{-1} at 2 GHz, assuming a 100 ms integration time and scan time of 60 s. This is easily achieved.

3.9. Digital Back End

Thirty-eight ROACH-2 systems process bands A and B from the 19 on-sky pixels. Some modifications were made to the chassis so that all connections and buttons, including power, are on the same end of the chassis, and a splitter was installed to feed both I and Q channels with the sampling clock. These chassis were mounted vertically with 13 on each of three shelves in the digital rack. Sheets of conductive foam between the chassis cushion them from telescope vibrations and provide additional RFI protection.

IF levels going into the ROACHes are set to give an rms count at the 8 bit ADC (e2v semiconductors, EV8AQ160) of ~ 42 with the calibration vane inserted into the beam. On the sky, the power drops by almost an order of magnitude so that the rms is ~ 14 , allowing a significant margin above the digitization noise level. A Valon 5009 synthesizer generates the 2 GHz sampling clock, and the ADCs trigger on rising and falling edges to sample the input at 4 GHz.

Data collection is initiated by sending a signal to the ROACHes to start integrating on the next hardware 1 PPS (pulse per second) signal. A total of 8192 samples are then captured from each ADC and transformed to the frequency domain by a PFB (Price 2021) to produce 1024 complex frequency channels spaced by 1.953125 MHz.

Weighted combinations of the complex I and Q values yield the lower and upper sidebands. After squaring the channel amplitudes, 39,063 power spectra are accumulated for a 20.000256 ms integration time. Consecutive integrations follow on without any dead time. Final requantization of each integration reduces the value to a 24 bit integer without risk of overflow, and the LSB and USB spectra are transmitted with a sequence number to the host computer through a 10 Gb switch using UDP.

3.9.1. Polyphase Filter Bank Implementation

The four-tap PFB (see, e.g., Price 2021) takes 8192 voltage samples and applies weights given by the product of a




Figure 15. Effective channel passbands from the PFB.

Hamming window and a sinc function:

$$w_n = \left(0.46164 \cos \left(2\pi \frac{n - \frac{N+1}{2}}{N} \right) + 0.53836 \right) \cdot \text{sinc} \left(4\pi \frac{n - \frac{N+1}{2}}{N} \right). \quad (2)$$

Figure 15 shows the resulting passband. Although the standard calculation of the noise equivalent bandwidth indicates a sensitivity loss of $\sim 10\%$, simulations predict a value of only 2.5%. This is perhaps related to the nature of the PFB, in which consecutive spectra overlap by 75%.

3.9.2. IQ Separation and Calibration

To achieve a high degree of fidelity in the LSB and USB spectra, a set of IQ coefficients is used to combine the I and Q signals (Finger et al. 2013). Initially, these coefficients are set so that there is no separation of I and Q spectra. A tone from a Hittite HMC-2220 frequency synthesizer is injected at the center of each frequency channel, and the amplitudes and phase difference of I and Q are measured. These are used to compute the IQ coefficients, which are then written to the field-programmable gate array (FPGA). Thereafter, when the coefficients are used to weight the sums of I and Q, the resultants are the sideband-separated LSB and USB spectra.

As a practical detail in the calibration, the DCM1 IF amplifiers are turned off to remove most of the noise component. Injection of the test tone at a power level comparable to the normal noise power level exercises a similar range of bits in the ADC, but all the power is concentrated in a single channel of the PFB. With the normal scaling this would overflow the corresponding register, so a different rescaling is implemented relative to the normal scaling for noise.

3.9.3. Aliasing

Aliasing of signals outside the nominal spectrum occurs because the band-defining filters for bands A and B and the low-pass filters in DCM2 do not efficiently filter out signal




Figure 16. A Mini-Circuits low-pass filter (LPF), model LFCN-1700, is used to attenuate signals potentially aliased into the 0–2 GHz baseband into the EV8AQ160 ADC. The ADC response falls off slowly according to its data sheet. Currently, the sampling frequency folds in significant aliased power, but a proposed increase to 4.5 GHz will essentially eliminate this.

close to the band edge, as illustrated in Figure 16. In part this is because filters with very sharp cutoff were avoided since the high number of poles required makes them more temperature sensitive, potentially leading to subtle changes in the passband shape. Currently, channels where the contribution of unwanted signal is more than ~ -15 dB relative to the in-band signal are discarded. A better solution is to sample at 4.5 or even 4.25 GHz so that the aliased power is attenuated more by the filters, with most of the aliasing being folded to frequencies outside the band of interest. It is a significant challenge to meet the timing in the FPGA, so that was deferred until resources were available. That work is currently underway.

3.10. Monitor and Control

Much of the antenna-based monitor control system electronics were taken directly from CARMA (Woody et al. 2004), with some additional components based on a control system for Argus (Sieth et al. 2014). A VERSAmodule Eurocard crate with a CPU running Linux (CentOS 6.5, 32 bit) acts as a master in a controller-area-network bus (CANbus)¹⁸ distributed system, with a protocol developed for CARMA (Woody et al. 2007). CANbus modules include encoder and drive interfaces, environmental monitoring, calibration vane interface, cryogenic compressor monitoring, etc. Other parts of the system, such as the calibration signal synthesizer and the ROACH-2 modules, communicate directly with the central control system through the antenna 10 GbE switch.

Following the CARMA scheme, most monitor points are in a 0.5 s cadence “blinking frame.” COMAP places an additional requirement on the monitoring because the rapid scanning needs a faster update of the antenna pointing. In its fast motion modes, the antenna is unable to precisely track the commanded positions, but the actual position is always known accurately. Encoder readings are strobed on every 100 ms boundary (see Section 3.11 for details of timing) and are reported in sets of five per the 0.5 s packet. Simple interpolation onto the 20 ms

integration interval is more than accurate enough for mapping the data to position on the sky.

3.10.1. Warm Support Electronics

Custom warm support electronics control and monitor the focal plane array’s front-end receiver package with 38 independently settable bias voltages for the LNAs and monitors. A second system controls the 38 DCM2 second-stage downconverters, including their environmental monitor points and their LO phase-lock states. Two Netburner Mod5270 microcontrollers, one for the front end and one for the back end, communicate with the COMAP control system by Ethernet, and with the electronics through I2C serial commands. To suppress RFI, both microcontrollers are in individual shielded enclosures with media converters for fiber Ethernet connections and RF filtering on the 100 kHz I2C signal lines. Controlling individual components with I2C internal interfaces through sub-buses selected by cascaded I2C switches has proved to be a compact and practical solution to reaching the several hundred control and monitor points in the system.

We took several steps to reduce pickup and other noise in the front-end LNA bias system, minimizing the effects of connections to the bias electronics on overall system noise. Circuit cards for bias, I2C bus switches, LNA power control, and cryogenic diode temperature current sources and readouts are on a simple backplane housed in a small card cage mounted on the cryostat, connected by ~ 1 m long cables to the cryostat and by a ~ 10 m long cable to the control microcomputer in the receiver cabin. A 1:44.5 resistive voltage divider between the bias system and LNA gates attenuates pickup from outside the cryostat and limits the gate voltage to 0.3 V even if the bias voltage swings to an op-amp rail. Series diodes clamp the drain voltages to a safe 1.4 V, and the drain bias circuit limits the output current to 50 mA. Soft limits in the microcontroller firmware further ensure that the bias voltages on the gates and drains are within safe limits for the LNAs. Mechanical relays under microcontroller control can electrically isolate the bias electronics and LNAs and allow power sequencing under microcontroller control.

Additional electronics in the front-end system control and monitor the power supplies, source $10 \mu\text{A}$ current for the cryostat’s temperature-sensing diodes, connect with the cryostat’s pressure gauge, and measure temperatures within the electronics. Small cards in each DCM1 module control and monitor IF amplifier power and temperatures. The back-end system provides power and communication for each of the 38 DCM2 modules and also monitors the phase-lock bits for the downconverter LOs.

3.10.2. Control Software

The antenna control system is an amalgam of two preexisting systems. Code from the CARMA software system (Scott et al. 2004) that ran in the antenna Linux computer was modified to communicate with a user interface and archiving system based on code used in several other telescopes at OVRO and elsewhere (Padin et al. 2002; Story et al. 2012). The observer is presented with a flexible user interface that has a command window and user-definable monitoring windows, with either numerical fields or time series plots. For the most part, operations are run through a built-in scheduling language

¹⁸ <https://www.edn.com/canbus-networks-break-into-mainstream-use/>

that can also invoke Python scripts for easy development of functionality.

Monitor points from CAN modules in the antenna are pushed out every half second and collated in the CARMA-based antenna code. These are filled into registers in the central control system, where they are recorded in a flat file system. Additional housekeeping data, such as pointing-model-corrected azimuth and elevation, source name, and commanded position, are combined with the CANbus monitor points.

The control system also interrogates the warm support electronics for monitor data from the front end and the DCM2 modules on a 15 s cadence, and data from the site weather station are included in the registers.

Much of the antenna control code was taken directly from the CARMA system. Inputs from the main control system are at a high level since the antenna code can take input in the form of R.A. and decl., azimuth and elevation, and offset pointing in various coordinate systems. Conversion from J2000 coordinates to antenna azimuth and elevation, including refraction correction, is performed by the antenna computer. It also applies the pointing model, described below.

Leighton's original servo loop (Leighton 1977) is used to close the position loop. It takes the current and predicted positions at consecutive half-second intervals and derives velocities to apply to the azimuth and elevation motors. Proportional, integral, and differential terms are used to improve tracking accuracy.

Some special modes were implemented for COMAP, allowing the antenna to do horizontal raster scans or Lissajous figures on the sky. In these cases, the tolerance on the acquisition of the commanded position is relaxed and the telescope does not closely follow the nominal trajectory, but the position is recorded accurately every 100 ms. Raster scans are implemented by acquiring the source and then alternating positive and negative offsets, allowing the antenna to move between them at a speed determined by its natural acceleration and velocity limits. The Lissajous mode, where sinusoids with individual amplitudes, frequencies, and phases are applied to the azimuth and elevation axes, is less stressful than the raster scans, as well as giving better field coverage near transit. The relative merits for the data quality are analyzed in Foss et al. (2022).

3.11. Time and Frequency References

Universal Coordinated Time is derived from the OVRO Meinberg Lantime stratum 1 GPS time server. Computers on the OVRO network synchronize to this using the Network Time Protocol (NTP), typically maintaining time to better than a millisecond. A CARMA master clock CAN module takes the 10 MHz and 1 PPS outputs from the GPS clock and uses them to discipline its Rb clock (SRS PRS-10) with a 4 hr time constant, giving excellent short- and long-term stability. All systems requiring 10 MHz and 1 PPS inputs use the output from the Rb clock. Components that require the 10 MHz signal include the sampling clock for the digital data acquisition, the sampling clock frequency counter monitor, the first and second downconverter oscillators, and the Hittite calibration signal synthesizer.

Precise phase locking of all oscillators is critical to the system stability. Small changes in LO frequencies shift ripple slightly, producing residual ripple in the data.

Matching the time between the encoder readings and the data integrations is the most critical aspect of timing. All of the CAN modules in the system keep Modified Julian Date (MJD) time internally. Time stamps are sent out on the CANbus from the Linux host computer approximately every 10 s. The exact time is read from the NTP clock when the time message is transmitted, and since the time stamp messages have the highest priority on the bus, there is negligible latency, ensuring that timing is accurate to about a millisecond. Internally, the CAN module servos its timer to keep its clock synchronized without introducing discontinuities. Encoder readings are triggered by an interrupt from its MJD clock every 100 ms, and after five readings are acquired, they are sent back in the blanking frame packet on the 0.5 s boundary. The time assigned to each encoder reading is its nominal exact 100 ms time boundary, so latency in returning packets does not affect the accuracy. Note that this assumes that timing on the host is tight enough to avoid ambiguities in identifying the correct 0.5 s frame, which has been verified over a decade of operation in CARMA.

Accurate timing for the data integrations relies on the 1 PPS from the master clock, which is fed to each ROACH-2 chassis. When data acquisition is initiated from the control system, a "dada" ring buffer (van Straten et al. 2013) is set up to capture UDP data packets on the 10 GbE network. The header is set up, and in the half second before the designated start time a command is sent to the ROACH-2 PPCs to arm them to start acquisition on the next 1 PPS rising clock edge. Although the ROACH-2s are not guaranteed to start on the same 4 GHz sampling clock cycle, temporal offsets much smaller than the integration time have negligible effects.

Verification of the precise alignment of the CAN module time with the PPS was done by temporarily injecting a pulse generated by the CAN module software into one of the feed ROACH-2 ADC inputs and verifying that it appeared at the correct time in the data stream. This has a modulo 1 s ambiguity, which does not detect the case where a ROACH-2 misses its trigger time and is delayed by a second. That case can be detected by comparing the times that the spectrometer total power changes for all feeds as the calibration vane is inserted and removed. Affected data are flagged and removed from the analysis; this is infrequent enough that it is not deemed useful to try to realign the data after the fact.

3.11.1. Pointing

Pointing of the 10.4 m telescope has been extremely well characterized over decades of operation. An optical telescope is permanently mounted on a representative part of the backing structure and observes through a hole in the primary. A series of stars are observed over the sky, with the observer centering each in the field of view and recording the necessary offsets. An eight-parameter model is fit to these data. Five of these refer to the mount and are common to the radio pointing. Three (elevation, cross-elevation collimation and sag) must be separately measured for the radio beam through radio pointing. Sag is the most difficult term to determine for radio pointing, but once determined from a dedicated run over a large elevation range, it is stable enough not to require recalibration.

Tilting of the antenna azimuth axis driven by solar heating is significant, but a tiltmeter located on-axis continuously monitors the two orthogonal tilts, and the values are applied to the pointing model in real time. At intervals of a week or so,

the telescope is rotated through 360° in azimuth and the corotating tiltmeter readings analyzed to derive a constant offset between the tiltmeter axis and the azimuth axis. This is particularly important before a pointing run.

In general, blind pointing is accurate to $\leq 0'.2$, and offset pointing to $\leq 0'.02$, more than sufficient for COMAP science.

3.12. Data Acquisition and Processing

When data acquisition is started by the control system, a software package, PSRDADA (van Straten et al. 2013), is initialized with header information including such details as the start time, integration time, and other necessary metadata. Ring buffers are initialized to capture spectrometer data pushed out by the ROACH-2s as 24 bit integers, and they are combined into 20 ms frames. If packets are not captured within a specified time, the data are filled in with a special value. Data are streamed into a “dada” file. Typically, observations are split into 1 hr chunks to create manageably sized data files.

Once a dada file is complete, it is sent to campus to process into a “Level 1” file. This conversion includes transposing the ordering of feed and time, converting spectra from 24 bit integer to 32 bit floats, replacing “special” 24 bit numbers (e.g., overflow) to “NaN,” including monitor point archive values, calculation of R.A. and decl. from azimuth and elevation for each feed, and calculation of system temperatures.

The monitor point archive includes “feature bits” that are set by the control system to indicate particular observing modes. For example, they may indicate whether raster scanning or Lissajous scanning is in effect, or if the measurement is a sky dip or a system temperature measurement. The Level 1 conversion uses these bits to group or process data in particular ways, and they are also available for later stages of processing.

4. Performance

4.1. RF Performance

System noise performance was evaluated by scanning the telescope in elevation on the sky, bracketed by an ambient load measurement at the start and end. For each 20 ms integration, the system temperature and gain were determined for every 2 MHz channel using the ambient vane calibration technique. This system equivalent noise temperature, which is referred to outside the atmosphere, can be written as

$$T_{\text{sys}} = (T_{\text{rx}} + (1 - \eta_{\text{sp}})T_{\text{amb}} + \eta_{\text{sp}}(1 - e^{-A\tau_z})T_{\text{atm}})\eta_{\text{sp}}^{-1}e^{A\tau_z} + T_{\text{cmb}}, \quad (3)$$

where T_{rx} is the receiver noise temperature, η_{sp} is efficiency due to spillover terminated at ambient temperature (T_{amb}), τ_z is the zenith opacity, $A = 1/\sin\theta_{\text{el}}$ is the number of air masses as a function of the elevation angle (θ_{el}), T_{atm} is the physical temperature of the atmosphere, and T_{cmb} is the brightness temperature of the cosmic microwave background.

By fitting the data to Equation (3), the system parameters can be determined. However, there is sufficient degeneracy among some of the parameters that we fix $T_{\text{atm}} = T_{\text{amb}} - 20$ K and $\eta_{\text{sp}} = 97.5\%$. The spillover is consistent with a measurement made when the secondary support was removed for modification. System temperatures measured with and without the secondary assembly in place yielded an estimate of the spillover noise at zenith. For the two feeds installed in the initial configuration, the spillover was found to be $\sim 5\text{--}6$ K (Figure 17). In reality, the




Figure 17. Zenith spillover noise estimated by taking the difference in system temperature with the secondary support absent and installed.

spillover is a function of elevation, so it is partially degenerate with the atmospheric elevation dependence. For now we assume for simplicity that the spillover is a fixed number.

Noise temperatures derived for all the receiver channels are shown in Figure 18. In general, these are excellent, but an obvious issue is the very prominent noise spikes, particularly in the lower half of the band. The precise cause of these spikes is currently unclear, but it is known to be an interaction between the polarizers and the corrugated horns. When the polarizer is replaced by a plain circular waveguide, or the corrugated horn is replaced by a smooth-wall conical horn, the spikes disappear. Return loss measurements also show some evidence of the spikes, but electromagnetic simulations with HFSS with some tolerance errors introduced have so far failed to show similar behavior. An experiment with an attenuating septum in the circular waveguide between the rectangular-to-circular converter and the polarizer to terminate the orthogonal linear polarization resulting from reflections did not show any improvement. Similar phenomena have been reported by, for example, Morgan & Pan (2013), which will guide a resolution of the problem in future instruments.

In the interest of expediency, the receivers were installed in this condition, and the spikes are flagged out from the data. They are fairly stable in frequency and amplitude, but they have been observed to vary slightly with bias changes in the first stage of the LNA.

Figure 19 shows the gain of the receivers represented by the median over all the feeds. Spikes due to sampling artifacts have been removed, but there is a significant amount of structure remaining. Some of the uncertainty in the noise and gain measurements results from the calibration load variability. It was understood at the start that it would be almost impossible to build a calibrator to precisely determine the system passband. Finally, the derived atmospheric contribution is presented in Figure 20.

The intention was to use the observation scan data to remove structure in the passband and noise temperature and to use the calibration vane to obtain the overall scaling to brightness temperature units. Differencing ambient load measurements taken 1 minute apart, at 40° and 60° elevation, revealed a ripple component corresponding to the LNA-load distance. The temperature was stable to better than the resolution of the temperature sensor, 0.2 K, or 0.07%, and the mean difference in power for the two measurements was between -0.42% and 0.12% , depending on the feed. However, the standard deviation of the ripple varied among feeds from 0.23% to 1.49%. This has little impact on the CO line mapping results or the Galactic survey.

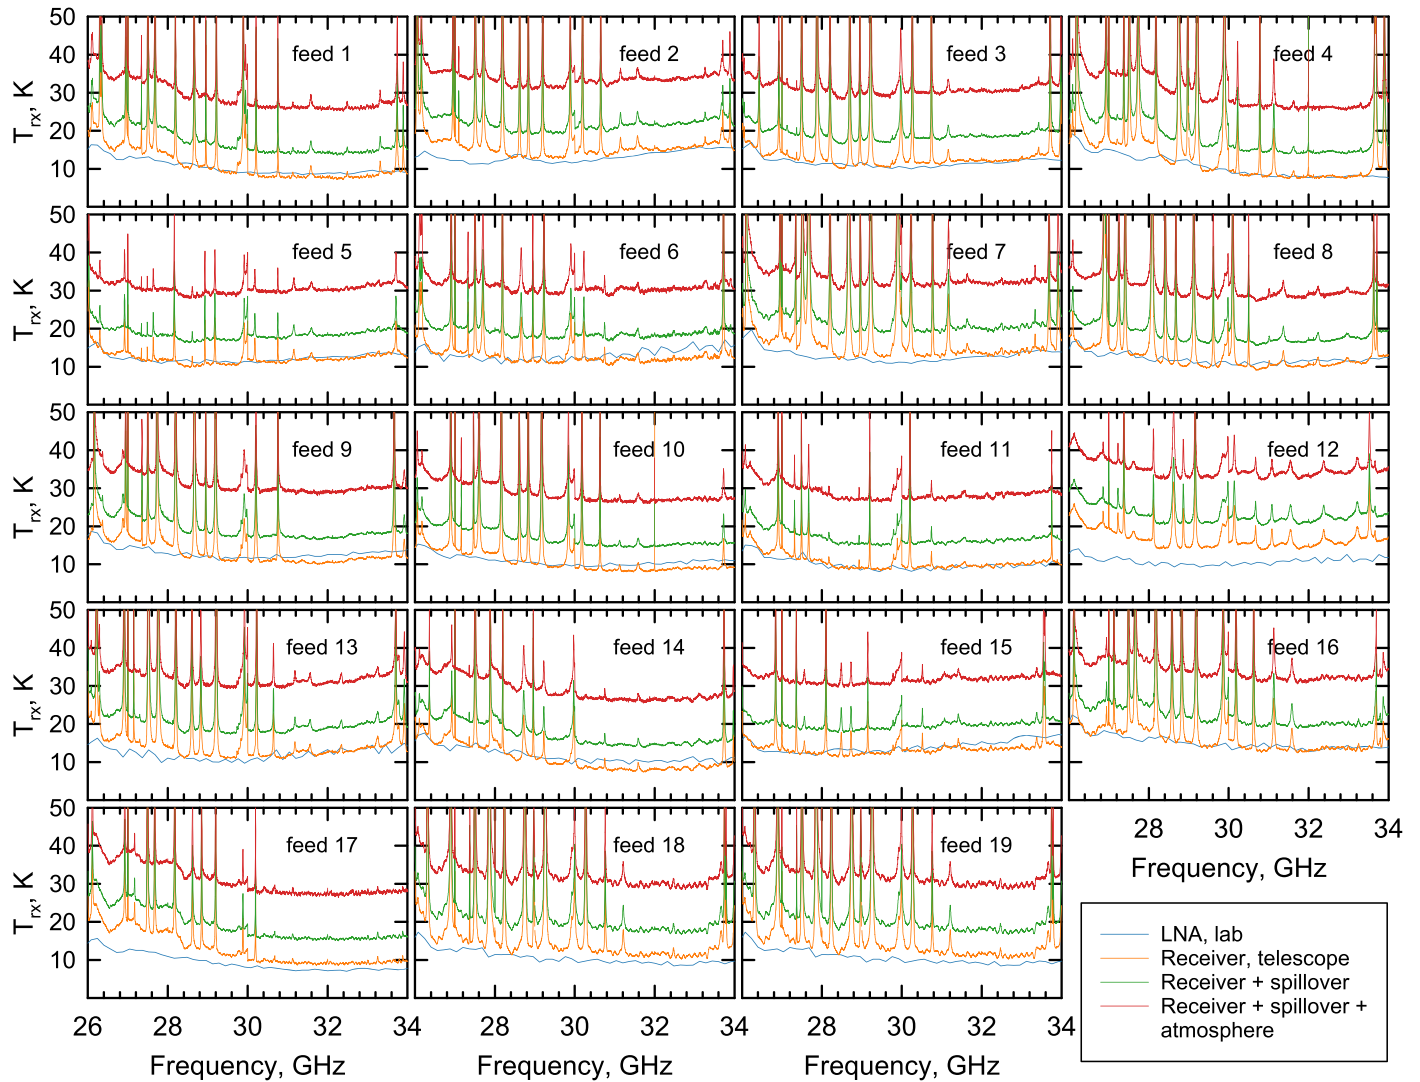


Figure 18. System noise temperatures. The blue curves are the LNA noise temperatures measured in the lab. The orange curves show the values derived from the sky dip, while the green curves also include the estimated spillover to ambient. Curves in red include all noise components, including the zenith atmospheric brightness component. See text for discussion of the noise spikes.

4.2. Spectrometer Performance

The analog signal chain is set up with the calibration vane inserted to give a noise input to the ADCs such that the $\pm 3\sigma$ levels are at the limits (-128 to $+127$ digitizer counts) so that $\sigma \approx 42$ counts. On the sky the power is about 10 dB lower, so $\sigma \approx 14$ and the digitization noise is negligible. At this setting the digitizer is accurate on both the sky and the calibration load and has a sufficient dynamic range to cope with the ~ 10 dB power variation across the spectra.

A typical raw spectrum is shown in Figure 21. The end channels of each sideband, corresponding to DC and its Nyquist alias, are discarded. The discontinuities at the center of the sidebands are artifacts of the four-core architecture of the ADCs, while the discontinuities at one-quarter and three-quarters of the channel range are attributed to the DC component of the 0–2 GHz baseband signal, or the 6 and 8 GHz LO signal at the IF. These frequencies are also discarded.

During observing, the IQ coefficients for separating the sidebands of the DCM2 LOs are periodically determined. This is essential when the ROACH-2s are power-cycled, but in general the coefficients are very stable over time. Power-cycling the

ROACH-2s causes them to lose clock alignment, and an arbitrary shift of up to 16 integer clock cycles can be introduced. In principle, only a coarse sampling of 16 frequencies is sufficient to determine this, and the previous coefficients can be modified accordingly, but since this is required infrequently, it has not been worth the time to implement.

Figure 22 shows the achieved rejection ratio after calibration for a typical band. The 20 dB specification is easily met at all frequencies of interest. Furthermore, the calibration is stable over the longest period that has been observed between calibrations (of order a month).

4.2.1. Aliasing

Aliasing was a known artifact of the system during the design phase. Sampling at 4 GHz was intended as an interim solution, with a goal of 4.25 GHz sampling (Section 3.9.3). Tests showed that the current FPGA implementation does not support sampling more than a few percent higher than 4 GHz, so the aliasing was quantified in detail using the injected test tone.

As expected, there is significant leakage from signals outside the band being reflected into the band. Contamination from just




Figure 19. Median gain for the 19 receiver channels.




Figure 20. Atmospheric zenith opacity derived from an elevation scan on the sky on a good observing day. The ensemble of feeds is shown in gray, while the dashed and solid lines are modeled with precipitable water vapor of 6 and 9 mm, respectively.


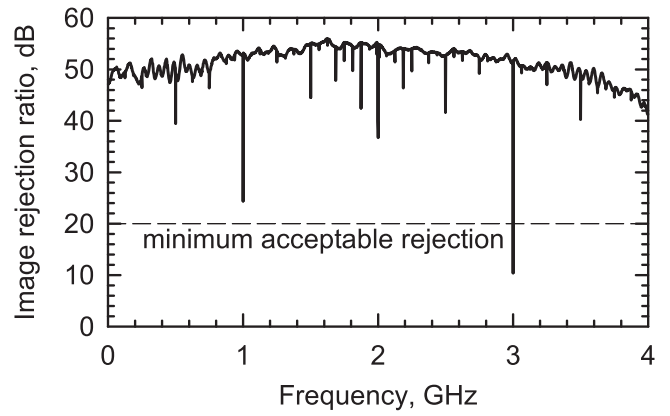


Figure 21. Example of a raw spectrum. Dashed lines indicate boundaries of the four sidebands.


outside the opposite edge of the band was also found in the measured data (see Figure 23). This results from aliasing into the opposite sideband that is not rejected by the IQ separation procedure, as illustrated in Figure 24. The IQ coefficients determined for the direct signal do not result in good separation for the aliased signal, though this can potentially be improved with better analog path equalization.

4.3. System Stability

By design, the COMAP system should be very stable over timescales of minutes apart from fundamental known physical effects, particularly “ $1/f$ ” amplifier gain fluctuations (Escotte & Goneau 2013; Weinreb & Schlee 2014) and atmospheric



(a)



(b)

Figure 22. (a) A typical plot of the image rejection ratio across the two 2 GHz sidebands of a second IF. The requirement is 20 dB, which is easily met except at frequencies that are rejected for other reasons. (b) The evolution of the median rejection ratio for all the bands in the system over a 1 month period. Arrows indicate time that the IQ coefficients were remeasured, usually because of a system restart.

opacity variations. However, several other effects were exposed during commissioning and routine observing.

Some of the IF amplifiers used proved to have unreliable bonding to the connectors, and significant changes in the signal amplitude as a function of time were noted. These were more significant at the lower frequencies consistent with variable capacitive coupling of a broken connection. Replacing affected amplifiers resolved this issue.

Some of the LNAs exhibited changes in gain under certain circumstances. One of the symptoms was a change in the bias current when the calibration vane was inserted. This was ascribed to a change in power at the second stage that was rectified in the gate, changing the bias voltage. Choosing a different gate bias was sufficient to eliminate this. Another change in bias was associated with the pointing direction of the antenna and hypothesized to be due to light affecting the first stage of the LNA. Testing suggested that both visible light and infrared light were implicated. A thin black polyethylene sheet was added behind the foam weather window, and an expanded polystyrene disk was inserted between the feed horns and the polyethylene to attenuate the infrared.

Ripples in the spectra were sometimes seen on windy days at certain angles of incidence of the wind on the receiver. The




Figure 23. Measured aliasing using an injected tone. The top panel shows the leakage from signals in the same sideband as the nominal signal, while the bottom panel shows signals leaking from the opposite sideband. The IQ coefficients applied are correct for the nominal signal but should not be expected to be appropriate for the image, which is 4 GHz away. The I and Q path lengths in DCM2 were not equalized, which would probably have given better cancellation of the aliased signal also.

period corresponded to a standing wave between the weather window and the LNA input. The polystyrene foam disk mentioned above was modified to have a slightly convex outer surface that would put the weather window under tension and stabilize it. Note that it is the stability of the standing wave that matters, rather than merely its presence.

During early tests, variations in the power were measured in the spectrometer without corresponding changes in the powers monitored in the DCM2s. This appeared to be correlated with temperature changes in the digital rack, and lab tests confirmed that the ADC gains are temperature sensitive. The measurement indicated a gain sensitivity of $\sim -0.34\%/^{\circ}\text{C}$, roughly twice the specification in the data sheet. Better temperature regulation was obtained after maintenance work on the air conditioning unit and relocation of its temperature sensor.

The Allan variance (Allan 1966) of the time series for each channel from the spectrometer quantifies how long an integration is consistent with statistical noise. For a series of consecutive channel power integrations, each of length τ , the Allan variance is

$$\sigma_{\text{av}}^2 = \frac{1}{2} \left\langle (\bar{p}_{n+1} - \bar{p}_n)^2 \right\rangle. \quad (4)$$

In an ideal statistical-noise-limited system,

$$\sigma_{\text{av}}^2 = \frac{1}{\Delta f \tau}, \quad (5)$$

but generally, as τ tends to longer times, the Allan variance will assume larger values owing to, for example, $1/f$ noise. If τ is




Figure 24. Diagram illustrating the effects of signal aliasing. The top panel shows the input from IF 1 to the second downconverter. The dashed box indicates the nominal band. In the second panel, the signal has been converted down to the 0–2 GHz baseband (IF2), folding the USB and LSB bands together, along with the aliased signals. Finally, in the bottom panel the digitally separated signals are shown with their assigned frequencies. Imperfect sideband separation is represented by the bottom hatched region, and the aliased signals are depicted by the arrows.

less than this inflection point, integrations may be summed to give close to theoretical noise.

Figure 25 shows the Allan variance on a range of timescales for all channels of one spectrometer. In the top panel of Figure 25 no processing has been applied to the data, and it is apparent that it is not possible to integrate more than a few seconds before the sensitivity starts being compromised. The plot contains no information on how correlated the channels are, but subtracting the mean value of all channels at each time step, any common mode is removed. The bottom panel of Figure 25 shows that doing this pushes out the usable integration time up by about an order of magnitude. The common mode is an expected consequence of the gain fluctuations of the LNAs (quantified by a single parameter given as the product of the transistor mutual




Figure 25. Allan variance normalized to $1/\Delta f \tau$ for the data from feed 1, band A, USB for a 10-minute stare at zenith. Top panel: values for the raw data stream. Bottom panel: variance calculated after removing a common-mode value from all channels.

gains), and the atmospheric brightness changes. The gain fluctuations are common to all four bands, while the atmospheric contributions have a significant commonality among feeds.

This very simple data processing indicates that scans over the source on timescales of ~ 20 s can potentially result in statistical-noise-limited measurements. Results presented in the data pipeline and science results companion papers use much more sophisticated procedures that mitigate a variety of systematic effects.

4.4. Interference

Generally RFI from coherent sources such as satellites and radio links is not a significant problem at the observing

frequency of COMAP. There are relatively few such sources, and the high directivity of the telescope minimizes pickup. Incoherent sources, such as power line arcing and ignition noise, are too weak to be detected at such high frequencies. Certain models of cell phones were found to produce noticeable effects when operated in the sidecab, but this is easily avoided.

During operation, certain broadband spikes were encountered in the data at certain times of year. These were found to be due to birds flying in front of the telescope aperture, or in front of the feeds. A different signature was seen from insects crossing the feed array, and the time resolution of the data enabled mapping their flight path over the pixels. Both the bird and insect signals vary significantly from year to year and with season and were infrequent and obvious enough to flag with custom spike detection code.

The Sun and the Moon were also found to contribute fluctuations to the detected power. As described in Section 3.1, there are some wide-angle sidelobes resulting from scattering by the secondary support struts. If the Sun or the Moon is present in those sidelobes, the data are compromised, so data are flagged when either is closer to the main beam than about 60° .

4.5. Verification on Astronomical Sources

4.5.1. Jupiter Observations

Regular observations are made of Jupiter to verify performance and check pointing. These are constant elevation scans that require different processing than the CO fields that suppress continuum sources. Representative beam patterns (Figure 26) agree well with predictions, showing the modest change in beamwidth with frequency.

Pointing and beam shapes are confirmed for all feeds. Figure 27 plots the beams in the nominal coordinates predicted for their pointing centers. Deviations from those are at the limit of the measurement accuracy.

4.5.2. Beam Maps

Verification of the beam shape and offset using Jupiter as an unresolved source are in good accord with the predictions. The maps for all the feeds are shown in Figure 27.

4.5.3. Spectral Line Observations

A project to map significant portions of the Galactic plane has yielded significant detections of several $H\alpha$ radio recombination lines (Figure 28) and potential detections of C and He recombination lines (Rennie et al. 2022). These serve as an important confirmation of the mapping between ROACH-2 channels and sky frequency. Additionally, they demonstrate that the instrumental spectral baselines are well enough controlled to produce high-significance line measurements.

5. Discussion

The COMAP Pathfinder was conceived as a proof of concept for a line intensity instrument that would allow exploration of the challenges posed by deep integration with a total-power system. It was also expected to place constraints on the range of models of CO distribution at redshifts $z = 2.4$ – 3.4 . The rapid deployment was made possible by the availability of a high-performance 10.4 m antenna. With a modification of the optics,




Figure 26. Measured and computed beam patterns for feed 1. The dots represent the measurements of the beam over all azimuthal angles around the boresight direction. Data are averaged over 1 GHz bandwidth. Solid lines show the computed patterns at azimuth angles of 0°, 45°, and 90°.




Figure 27. Beam maps derived for each feed from raster scans on Jupiter. Note that these are reversed left to right and top to bottom compared to a map of Jupiter since, e.g., pointing above Jupiter is a measure of a point in the beam below boresight.

including a new secondary mirror, it was well suited to a 19-feed 26–34 GHz receiver.

Design efforts were focused on the stability of the gain, particularly as a function of frequency, on timescales of a couple of tens of seconds. Amplifier gain instabilities are tied to the fundamental physics of the transistors, but since they appear as a common mode across the band, they can be filtered out. Similarly, fluctuations of the atmospheric brightness are very smooth in frequency and have a strong common-mode component among the feeds. Most of the problematic changes are related to standing waves in transmission paths. The data analysis is insensitive to ripple due to standing waves but very sensitive to changes in the ripple, typically caused by temperature changes or mechanical stress from the antenna motion.




Figure 28. Radio recombination lines observed in the COMAP passband in the star-forming region W43. The vertical dashed line indicates the expected line center due to the known velocity of the region.

Standing waves in the optics are particularly troublesome since they are at the highest frequencies. By using circular polarization, standing waves are greatly reduced. During the first observing season, two feeds had no polarizers, two had single-section polarizers, and the remainder had two-section polarizers. An unanticipated consequence was that interactions between the polarizers and corrugated feeds caused strong noise spikes in the band. However, the overall improvement in the data quality from the analysis pipeline with the polarizers installed outweighed the loss in discarding the affected frequency channels. All feeds now have the two-stage polarizers. Further development will include eliminating the spikes while retaining the polarizers.

Ripple due to a standing wave between the LNAs and the weather window was found to cause problems under certain conditions of wind and antenna orientation. This was easily solved by adding a foam backing to hold the window taut.

With careful selection of cable types, good temperature regulation, and avoidance of mechanical stress, cable standing waves proved not to be a limitation. In certain instances where cable standing waves were identified, it was traced back to a break in a connector or component, which was fixed or replaced.

The most concerning limitation of the system was revealed to be pickup from ground spillover and bright astronomical sources. While this was anticipated, it was not easy to quantify in advance. With appropriate scanning patterns and avoiding bright sources such as the Sun, good data are still achievable. However, for future experiments much of this type of contamination can be avoided with a dual-offset antenna. Mechanically, offset antennas are more massive to achieve the same stiffness and have a larger moment of inertia and therefore higher drive requirements. For the Pathfinder, the choice of an existing, well-understood, high-performance antenna was a critical factor in getting good-quality data on a short timescale that significantly outweighed potential benefits of a new unevaluated offset antenna.

Key requirements for the instrument have been informed by continuing theoretical simulations (Breyse et al. 2022; Chung et al. 2022), while the pipeline analysis (Foss et al. 2022) and science data processing (Ihle et al. 2022; Rennie et al. 2022) have been crucial in evaluating performance, detecting

problems, and verifying fixes. Since the first season of observations, several improvements have been made, such as upgrading LNA modules, installing two-stage polarizers on all feeds, and repairing faulty components. So far, no fundamental limits to this method of LIM have been encountered in the first season of observing, and the current and imminent modifications will improve the data quality for future seasons. Future instruments, such as a complementary 16 GHz system (Cleary et al. 2022), will benefit greatly from the experience with the COMAP Pathfinder.

This material is based on work supported by the National Science Foundation under grant Nos. 1517108, 1517288, 1517598, 1518282, and 1910999 and by the Keck Institute for Space Studies under “The First Billion Years: A Technical Development Program for Spectral Line Observations.” Parts of the work were carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration, and funded through the internal Research and Technology Development program. D.T.C. is supported by a CITA/Dunlap Institute postdoctoral fellowship. The Dunlap Institute is funded through an endowment established by the David Dunlap family and the University of Toronto. C.D. and S.H. acknowledge support from an STFC Consolidated grant (ST/P000649/1). J.B., H.K. E., M.K.F., H.T.I., J.G.S.L., M.R., N.-O.S., D.W., and I.K.W. acknowledge support from the Research Council of Norway through grants 251328 and 274990 and from the European Research Council (ERC) under the Horizon 2020 Research and Innovation Program (grant agreement No. 819478, COSMO-GLOBE). J.G. acknowledges support from the University of Miami and is grateful to Hugh Medrano for assistance with cryostat design. L.K. was supported by the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 885990. J. Kim is supported by a Robert A. Millikan Fellowship from Caltech. At JPL, we are grateful to Mary Soria for assembly work on the amplifier modules and to Jose Velasco, Ezra Long, and Jim Bowen for the use of their amplifier test facilities. H.P. acknowledges support from the Swiss National Science Foundation through Ambizione grant PZ00P2_179934. P.C. B. is supported by the James Arthur Postdoctoral Fellowship. R.R. acknowledges support from ANID-FONDECYT grant 1181620. M.V. acknowledges support from the Kavli Institute for Particle Astrophysics and Cosmology. We thank Isu Ravi for her contributions to the warm electronics and antenna drive characterization. We would also like to acknowledge the important contributions of Russell Keeney, Michael Virgin, and Andres Rizo for constructing and installing major mechanical parts of the instrument. We gratefully acknowledge hardware donations from Xilinx within the Xilinx University Program. Finally, we thank the anonymous referee for the careful reading of the manuscript.

Appendix A Design Principles

As a total-power instrument, COMAP must discriminate a tiny signal from a system noise level that is about six orders of magnitude greater. This system noise results from the instrument (antenna and receiver), the environment (ground and atmosphere, radio interference), and foreground astronomical sources (diffuse and compact). The key goal was to build a

system that is stable over times comparable to or longer than the time to scan a source. Faster scanning allows the source to be covered in a shorter time, but particularly for an existing telescope there is a limit to the maximum rate. For the 10.4 m Leighton antenna, the typical time to scan over a degree-sized patch is a few seconds. Ideally, fluctuations in the total power should be limited by the statistics of the radiometer (statistical) noise. In practice, there are some background fluctuations that are common to all frequencies or pixels that can be easily removed, but other subtler artifacts can easily dominate the data.

Most phenomena fall into one of two categories. These are essentially multiplicative (gain and passband changes) or additive ground pickup, bright astronomical sources. For COMAP, these are dominated by standing waves and sidelobes, respectively.

A.1. Standing Waves

Frequency-dependent variations in the output spectra of the instrument are many orders of magnitude greater than the deviations due to the expected CO signal and cannot be determined a priori to an acceptable accuracy ($\lesssim 10^{-7}$). The structure is primarily due to the passband shape of the receiver and optics, and to a lesser extent to variations of noise temperature with frequency. Details of how the data analysis deals with distinguishing the spatial and temporal structure of the source from the instrument and foregrounds are discussed by Foss et al. (2022). For the purposes of this paper, we can summarize the goal as reducing the temporal systematic variations in the spectrometer output to a level comparable to or less than the statistical noise over a time of order 10–30 s after simple temporal and frequency slopes have been removed. Provided that the systematic residuals are random among different observations and feeds, they will not add coherently and will remain low compared to the statistical noise.

Ripple in transmission paths is a dominant source of spectral structure, so stabilizing transmission path lengths was a major focus in the instrument design. Standing waves are produced by discontinuities between sections of a transmission path. Multiple standing waves generate periodic ripple in the passband of a system, with a periodicity related to the time delay of wave transmission between pairs of discontinuities. Any pair of discontinuities give rise to ripple that will correspond to a particular value in the k -spectrum.¹⁹ For a lossless transmission line with a length l and small voltage reflection coefficients r_1 and r_2 at the two ends, the complex amplitude transmission function is approximately

$$h(\omega) = (1 - r_1 r_2)(1 + r_1 r_2 e^{i\omega l/2nc}), \quad (\text{A1})$$

where ω is the angular frequency of the wave, c is the speed of light in vacuum, and n is the propagation velocity in the medium relative to vacuum. When applied to a white-noise spectrum with unit spectral density, the output spectrum has a

¹⁹ Note that a standing wave produced by a *single* discontinuity does not result in ripple in the passband.

sinusoidal ripple given by

$$\begin{aligned} P(\omega) &= P_r \cos \omega\tau, \\ P_r &= 2(1 - r_1 r_2)^2 r_1 r_2, \\ P_0 &= 2(1 - r_1 r_2)^2 (r_1^2 r_2^2 + 1), \\ \tau &= 2l/nc, \end{aligned} \quad (\text{A2})$$

where we assume without much loss of generality that the reflection coefficients are real. In the data processing every frequency channel is initially normalized against its running mean, so if this ripple in the power is constant in time, it does not contaminate the astrophysical power spectrum. Commonly, the reflection coefficients are relatively stable if proper connections are made, but the effective length of the line can vary on different timescales owing to temperature or flexing, for example, causing instability in the ripple pattern. Treating the reflection coefficients as small, and assuming an incremental change in the time delay of the line of $\delta\tau$, we arrive at a residual ripple in the power spectrum given by

$$\delta P(\omega) = P_r \delta\tau \omega \sin \omega\tau. \quad (\text{A3})$$

Assume that we take a series of spectra over a time t_{scan} and that τ changes linearly over that time by an amount $\delta\tau$. For each of the integrated spectra we can calculate the standard deviation due to the ripple across each spectrum and then compute the mean standard deviation over all spectra. For small changes in the line length and some reasonable approximations, the average standard deviation can be written as

$$\sigma_r = \frac{\pi}{\sqrt{2}} \delta\tau P_r f_{\text{mean}}, \quad (\text{A4})$$

where f_{mean} is the mean frequency across the spectrometer band.

It is evident that the effect is more significant for higher frequencies. Standing waves in the optics at 26–34 GHz are an order of magnitude more sensitive to length changes than are the standing waves in the cables to the spectrometers at 0–2 GHz.

Comparing the standard deviation of the ripple with the statistical noise aids understanding of the stability requirements. If each spectrum has a channel bandwidth Δf and is integrated for a time t_{int} , then the standard deviation of the power across the spectrum is

$$\sigma_{\text{stat}} = \frac{1}{\sqrt{\Delta f t_{\text{int}}}}. \quad (\text{A5})$$

If we require $\sigma_r \ll \sigma_{\text{stat}}$ over a scan time t_{scan} and the spectra are integrated over a time t_{int} , then the rate of change of the line length should ideally be limited to

$$\frac{d\tau}{dt} \ll \frac{2}{\pi f_{\text{mean}} P_r t_{\text{scan}} t_{\text{int}}^{1/2} \Delta f^{1/2}}. \quad (\text{A6})$$

Typically, this criterion may be rewritten as a rate of temperature change, since temperature is the main driver of the line length, l . Quantitative examples are given in the main text.

Particular regard was paid to the optical standing wave because of the sensitivity at the high frequencies as noted above. The three main strategies applied to mitigate this were

use of circular polarization, implementation of a diverter cone at the secondary, and thermal compensation of the receiver support. Significant suppression of the optical standing wave ripple is achieved by using circular polarization as detailed in Section 3.5.3.
















In the rest of the system, LMR cable (Times Microwave Systems) was used in most of the signal paths for several reasons. The dielectric is polyethylene, which does not suffer from the “knee” in the delay vs. temperature curve around 20°C observed in the more common PTFE (Czuba & Sikora 2011), which would make the standing waves very susceptible to temperature changes. LMR cable (apart from the smallest diameter) uses foamed polyethylene, which reduces the loss, as well as reducing temperature sensitivity further. Any standing wave corresponds to a particular k -value. It was not possible to make all the cables short or long enough to place their ripple period out of the COMAP k -value range, though that was done where possible.

A.2. Antenna Spillover and Sidelobes








Another focus in the design was control of the spillover of the telescope. Radiation entering from the ground and strong celestial sources is a particular problem because it varies with the telescope pointing, making it difficult to separate from the target signal. Since the telescope chosen for COMAP is a symmetric Cassegrain, there are limits on scattering due to unavoidable blockage by the secondary and its support struts.

Any parts of the antenna radiation pattern²⁰ not covering the target field have the potential to add contaminating radiation from other sources. Primarily, these are the Sun, the Moon, and Earth. Suppression below the signal level requires the pattern integrated across the source to be more than 80 dB below the integrated main beam. For the COMAP there are sidelobes high enough that the Sun and Moon significantly exceed this, and this has to be taken account of in the observing strategy. Ground radiation, while having a lower brightness temperature than the Sun, covers a much greater solid angle and therefore is difficult to avoid. Primarily it enters through fields diffracted at the secondary mirror perimeter that spill past the primary. The optical design was targeted at minimizing ground spillover, as discussed in Section 3.2.

ORCID iDs

James W. Lamb  <https://orcid.org/0000-0002-5959-1285>
 Kieran A. Cleary  <https://orcid.org/0000-0002-8214-8265>
 Dongwoo T. Chung  <https://orcid.org/0000-0003-2618-6504>
 Stuart E. Harper  <https://orcid.org/0000-0001-7911-5553>
 Andrew I. Harris  <https://orcid.org/0000-0001-6159-9174>
 Håvard T. Ihle  <https://orcid.org/0000-0003-3420-7766>
 Timothy J. Pearson  <https://orcid.org/0000-0001-5213-6231>
 Liju Philip  <https://orcid.org/0000-0001-7612-2379>
 J. Richard Bond  <https://orcid.org/0000-0003-2358-9949>
 Patrick C. Breyse  <https://orcid.org/0000-0001-8382-5275>
 Delaney A. Dunne  <https://orcid.org/0000-0002-5223-8315>
 Hans Kristian Eriksen  <https://orcid.org/0000-0003-2332-5281>
 Marie Kristine Foss  <https://orcid.org/0000-0001-8896-3159>
 Junhan Kim  <https://orcid.org/0000-0002-4274-9373>
 Hamsa Padmanabhan  <https://orcid.org/0000-0002-8800-5740>

²⁰ For simplicity, we treat the antenna as a transmitter; reciprocity dictates that the receiving gain pattern is identical to the transmitting pattern.

Anthony C. S. Readhead  <https://orcid.org/0000-0001-9152-961X>
 Rodrigo Reeves  <https://orcid.org/0000-0001-5704-271X>
 Thomas J. Rennie  <https://orcid.org/0000-0002-1667-3897>
 Nils-Ole Stutzer  <https://orcid.org/0000-0001-5301-1377>
 Marco P. Viero  <https://orcid.org/0000-0003-0545-4872>
 Duncan J. Watts  <https://orcid.org/0000-0002-5437-6121>
 Ingunn Kathrine Wehus  <https://orcid.org/0000-0003-3821-7275>

References

- Allan, D. 1966, *IEEEP*, **54**, 221
 Born, M., & Wolf, E. 1980, *Principles of Optics* (6th ed.; Oxford: Pergamon)
 Breyse, P. C., Chung, D. T., Cleary, K. A., et al. 2022, *ApJ*, **933**, 188
 Carilli, C., & Walter, F. 2013, *ARA&A*, **51**, 105
 Chung, D. T., Breyse, P. C., Cleary, K. A., et al. 2022, *ApJ*, **933**, 186
 Clarricoats, P. J. B., & Olver, A. D. 1984, *Corrugated Horns for Microwave Antennas* (London: Peter Peregrinus)
 Cleary, K. A., Borowska, J., Breyse, P. C., et al. 2022, *ApJ*, **933**, 182
 Czuba, K., & Sikora, D. 2011, *AcPPA*, **119**, 553
 Emerson, D. T., & Payne, J. M. (ed.) 1995, in *ASP Conf. Ser. 75, Multi-Feed Systems for Radio Telescopes* (San Francisco, CA: ASP)
 Escotte, L., & Gonneau, E. 2013, in *22nd Int. Conf. on Noise and Fluctuations (ICNF)* (Piscataway, NJ: IEEE),
 Finger, R., Mena, P., Reyes, N., Rodriguez, R., & Bronfman, L. 2013, *PASP*, **125**, 263
 Foss, M. K., Ihle, H. T., Borowska, J., et al. 2022, *ApJ*, **933**, 184
 Helszajn, J. 1978, *Passive and Active Microwave Circuits* (New York: Wiley)
 Hoppe, D. 1993, *Scattering-matrix Program for Circular Waveguide Components*, NPO-18708, NASA Tech Briefs
 Ihle, H. T., Borowska, J., Cleary, K. A., et al. 2022, *ApJ*, **933**, 185
 Kerr, A. R., Wollack, E., & Horner, N. 1999, *Waveguide Flanges for ALMA Instrumentation* NRAO Tech. Rep. 278, <http://legacy.nrao.edu/alma/memos/html-memos/abstracts/abs278.html>
 Kovetz, E., Breyse, P. C., Lidz, A., et al. 2019, *BAAS*, **51**, 101
 Leighton, R. B. 1977, Final Technical Report for National Science Foundation Project AST 73-04708, *California Institute of Technology*
 Li, T. Y., Wechsler, R. H., Devaraj, K., & Church, S. E. 2016, *ApJ*, **817**, 169
 Maas, S. A. 1992, *Microwave Mixers* (2nd ed.; Boston, MA: Artech House)
 Morgan, M., & Pan, S. 2013, *ITST*, **3**, 72
 Padin, S., Shepherd, M., Cartwright, J., et al. 2002, *PASP*, **114**, 83
 Padman, R., & Hills, R. E. 1991, *IJIMW*, **12**, 589
 Pancharatnam, S. 1955, *InMS*, **41**, 137
 Parsons, A., Backer, D., Siemion, A., et al. 2008, *PASP*, **120**, 1207
 Penzias, A., & Burrus, C. 1973, *ARA&A*, **11**, 51
 Plambeck, R., & Engargiola, G. 2010, *Circular polarizers for the CARMA 1 mm receivers*, CARMA Memorandum Series #54, https://www.mmarray.org/memos/carma_memo54.pdf
 Price, D. C. 2021, *The WSPC Handbook of Astronomical Instrumentation* (Singapore: World Scientific), 159
 Rennie, T. J., Harper, S. E., Dickinson, C., et al. 2022, *ApJ*, **933**, 187
 Rudge, A. W. et al. (ed.) 1986, in *The Handbook of Antenna Design*, IEE Electromagnetic Wave Series, Vol. 1 and 2 (2nd ed.; Norwich: Peter Peregrinus)
 Rusch, W. V. T., Sørensen, O., & Baars, J. W. M. 1982, *ITAP*, **30**, 786
 Scott, S. L., Amarnath, N. S., Beard, A. D., et al. 2004, in *ASP Conf. Ser. 314, Astronomical Data Analysis Software and Systems (ADASS) XIII*, ed. F. Ochsenbein, M. G. Allen, & D. Egret (San Francisco, CA: ASP), 768
 Sieth, M., Devaraj, K., Voll, P., et al. 2014, *Proc. SPIE*, **9153**, 91530P
 Story, K., Leitch, E., Ade, P., et al. 2012, *Proc. SPIE*, **8451**, 84510T
 Tacconi, L. J., Genzel, R., & Sternberg, A. 2020, *ARA&A*, **58**, 157
 Tang, Y.-L., Wadefalk, N., Morgan, M. A., & Weinreb, S. 2006, in *IEEE MTT-S Int. Microwave Symp. Digest* (Piscataway, NJ: IEEE), 81
 Thomas, B. M. 1978, *ITAP*, **26**, 367
 Ulich, B. L., & Haas, R. W. 1976, *ApJS*, **30**, 247
 van Straten, W., Jameson, A., & Osłowski, S. 2013, *PSRDADA: Distributed Data Acquisition*, Astrophysics Source Code Library ascl:2110.003
 Weinreb, S., & Schlee, J. 2014, *ITMTT*, **62**, 83
 Woody, D. P., Beasley, A. J., Bolatto, A. D., et al. 2004, *Proc. SPIE*, **5498**, 30
 Woody, D. P., Wiitala, B., Scott, S. L., et al. 2007, *RSci*, **78**, 094501