The C-Band All-Sky Survey (C-BASS): Design and implementation of the northern receiver

O. G. King[1,2]*, Michael E. Jones[2], E. J. Blackhurst[3], C. Copley[2], R. J. Davis[3],
C. Dickinson[3], C. M. Holler[4,2], M. O. Irfan[3], J. J. John[2], J. P. Leahy[3], J. Leech[2],
S. J. C. Muchovej[1], T. J. Pearson[1], M. A. Stevenson[1], and Angela C. Taylor[2]

[1] *California Institute of Technology, Pasadena CA 91125, USA*
[2] *Sub-department of Astrophysics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford, OX1 3RH, UK*
[3] *Jodrell Bank Centre for Astrophysics, School of Physics & Astronomy, The University of Manchester,*
Oxford Road, Manchester, M13 9PL, UK
[4] *Hochschule Esslingen, Kanalstraße 33, Esslingen 73728, Germany*

ABSTRACT

The C-Band All-Sky Survey (C-BASS) is a project to map the full sky in total intensity and linear polarization at 5 GHz. The northern component of the survey uses a broadband single-frequency analogue receiver fitted to a 6.1-m telescope at the Owens Valley Radio Observatory in California, USA. The receiver architecture combines a continuous-comparison radiometer and a correlation polarimeter in a single receiver for stable simultaneous measurement of both total intensity and linear polarization, using custom-designed analogue receiver components. The continuous-comparison radiometer measures the temperature difference between the sky and temperature-stabilized cold electrical reference loads. A cryogenic front-end is used to minimize receiver noise, with a system temperature of $\approx 30\,\mathrm{K}$ in both linear polarization and total intensity. Custom cryogenic notch filters are used to counteract man-made radio frequency interference. The radiometer $1/f$ noise is dominated by atmospheric fluctuations, while the polarimeter achieves a $1/f$ noise knee frequency of 10 mHz, similar to the telescope azimuthal scan frequency.

Key words: instrumentation: polarimeters – Galaxy: general – radio continuum: general.

1 INTRODUCTION

The C-Band All-Sky Survey (C-BASS)[1] is an experiment to map the entire sky in total intensity and linear polarization at 5 GHz. C-BASS aims to provide high signal-to-noise ratio maps of the polarization of Galactic synchrotron radiation, largely uncorrupted by Faraday rotation. The survey is being conducted in two parts: for the northern sky, a new instrument has been deployed on a 6.1-m antenna at the Owens Valley Radio Observatory (OVRO), California, USA, while for the southern sky a similar instrument is used on a 7.6-m telescope at the MeerKAT (Booth & Jonas 2012) support site near Carnarvon, South Africa. The two systems are matched to each give the same resolution ($0\overset{\circ}{.}73$, Holler et al. 2012) and similar system temperatures of $\approx 30\,\mathrm{K}$, and the data will be merged to form a single set of all-sky images.

The primary purpose of the C-BASS experiment is to assist in the measurement of the polarized cosmic microwave background radiation (CMB) (King et al. 2010).

Measuring the B-mode polarization signal requires separation of the CMB from foreground emission, which at 5 GHz is dominated by diffuse Galactic synchrotron emission, to greater accuracy than is possible with our current understanding of Galactic foregrounds (Kogut et al. 2007; Gold et al. 2011; Macellari et al. 2011). C-BASS data are expected to improve the accuracy of foreground subtraction in, for instance, template fitting analyses (Ghosh et al. 2012).

The sensitivity requirement for the survey is set by requiring that there be a 5-sigma detection of polarization in at least 90% of the sky. Interpolating between existing maps at 1.41 GHz (Wolleben et al. 2006) and 23 GHz (Page et al. 2007) we obtained an estimate that 90% of the sky pixels would have a polarized intensity of 0.5 mK or greater, giving a sensitivity requirement of < 0.1 mK/beam in linear polarization. The instrument design will reach the same thermal noise level in intensity, although we expect the intensity maps to be confusion limited at a higher level, $\approx 1\,\mathrm{mK}$.

Achieving this sensitivity is a balance between bandwidth, system temperature, and integration time as dic-

* E-mail: ogk@astro.caltech.edu
[1] http://www.astro.caltech.edu/cbass/

tated by the radiometer equation (Kraus 1986):

$$\sigma = \frac{T_{\mathrm{sys}}}{\sqrt{\Delta\nu\tau}}, \qquad (1)$$

where σ is the noise level, T_{sys} is the system temperature, $\Delta\nu$ is the bandwidth, and τ is the integration time. While a broader bandwidth would make the target sensitivity easier to achieve, interference from man-made radio transmissions limited our effective bandwidth to $\approx 500\,\mathrm{MHz}$ in the northern survey, and a consequent thermal noise level of $1.3\,\mathrm{mK\,s^{-1/2}}$ given a 30 K system temperature. Reaching this noise level requires a low level of $1/f$ noise in the receiver: in a conventional radiometer the $1/f$ noise would exceed the thermal noise level set by the radiometer equation over the time-scale of the measurement. We overcome the $1/f$ noise problem by designing a receiver that suppresses $1/f$ noise, and using a destriping mapmaking algorithm (Sutton et al. 2009) to remove the $1/f$ noise that remains.

The C-BASS receiver is designed to measure both total intensity and linear polarization simultaneously, with minimal systematic error, in a single frequency channel between 4.5 and 5.5 GHz. To achieve this we have developed a hybrid receiver design in which the total intensity is measured in a continuous-comparison fashion by comparing the sky temperature against a cold electrical reference load, similar to the *Planck* LFI receivers (Davis et al. 2009). The linear polarization is measured by correlating orthogonal circular polarization signals from the horn. This architecture suppresses the instabilities due to $1/f$ noise in the receiver in both the intensity and polarization measurements. The C-BASS receiver is unique in its combination of a continuous-comparison radiometer and a correlation polarimeter in a single instrument. Some polarimeters, e.g. QUIET (Buder 2010; QUIET Collaboration 2012) and GEM (Bergano et al. 2011), measure polarization the same way that we do, but do not measure total intensity with the $1/f$ noise suppression of the continuous comparison architecture[2]. Similarly, some continuous-comparison receivers, e.g. *Planck* LFI (Davis et al. 2009), are able to measure one component of the linear polarization vector by differencing the power in orthogonal linear modes, but are only able to measure the full linear Stokes vector by rotating or changing the orientation of the receiver with respect to the sky. The hybrid architecture used in the C-BASS receiver does have several disadvantages: four receiver gain chains are needed instead of two, hence greater cost, and the addition of a $180°$ hybrid to the signal path before the first low-noise amplifier degrades the sensitivity of the instrument.

In this paper we describe the analogue receiver built for the northern survey. The northern and southern receivers have the same architecture and share identical cryogenic front-ends, but the southern receiver was built later and took advantage of advances in digital processing hardware to replace the rest of the receiver with real-time digital processing (C. Copley et al., in prep.). The project as a whole is described in the project paper (M.E. Jones et al., in prep.) and the on-sky commissioning of the northern instrument is described in S.J.C. Muchovej

et al., in prep. In Section 2 we describe the overall architecture of the system. Sections 3 and 4 describe in detail the implementations of the cryogenic receiver and warm electronics respectively. We describe the performance of individual sections of the receiver throughout the text, and in Section 5 present performance results for the full receiver. We begin each section with a qualitative outline, followed by a detailed description.

2 INSTRUMENT DESIGN

The northern C-BASS instrument is a combination of a continuous-comparison radiometer and a correlation polarimeter. A simplified schematic diagram of the instrument is shown in Fig. 1. The instrument consists of a cryogenic receiver followed by a warm radiometer/polarimeter and then a digital backend. In the cryogenic receiver, orthogonal linear polarizations are extracted from the sky signal by the orthomode transducer (OMT). These pass through a $90°$ hybrid (included in the OMT block in Fig. 1) to form orthogonal circular polarizations E_L and E_R. A calibration signal from a noise diode is weakly coupled in to each circular polarization using $-30\,\mathrm{dB}$ directional couplers. The circular polarization signals are then combined with independent reference signals $E_{\mathrm{ref\,1}}$ and $E_{\mathrm{ref\,2}}$ using $180°$ hybrids. The four signals thus formed, $(E_L \pm E_{\mathrm{ref\,1}})/\sqrt{2}$ and $(E_R \pm E_{\mathrm{ref\,2}})/\sqrt{2}$, are then amplified and filtered by independent, phase- and gain-matched, chains of components.

After further gain and band-pass definition filters, each signal is split in two, and sent to the radiometer and polarimeter branches. The radiometer and polarimeter functions are logically distinct, and are described separately below.
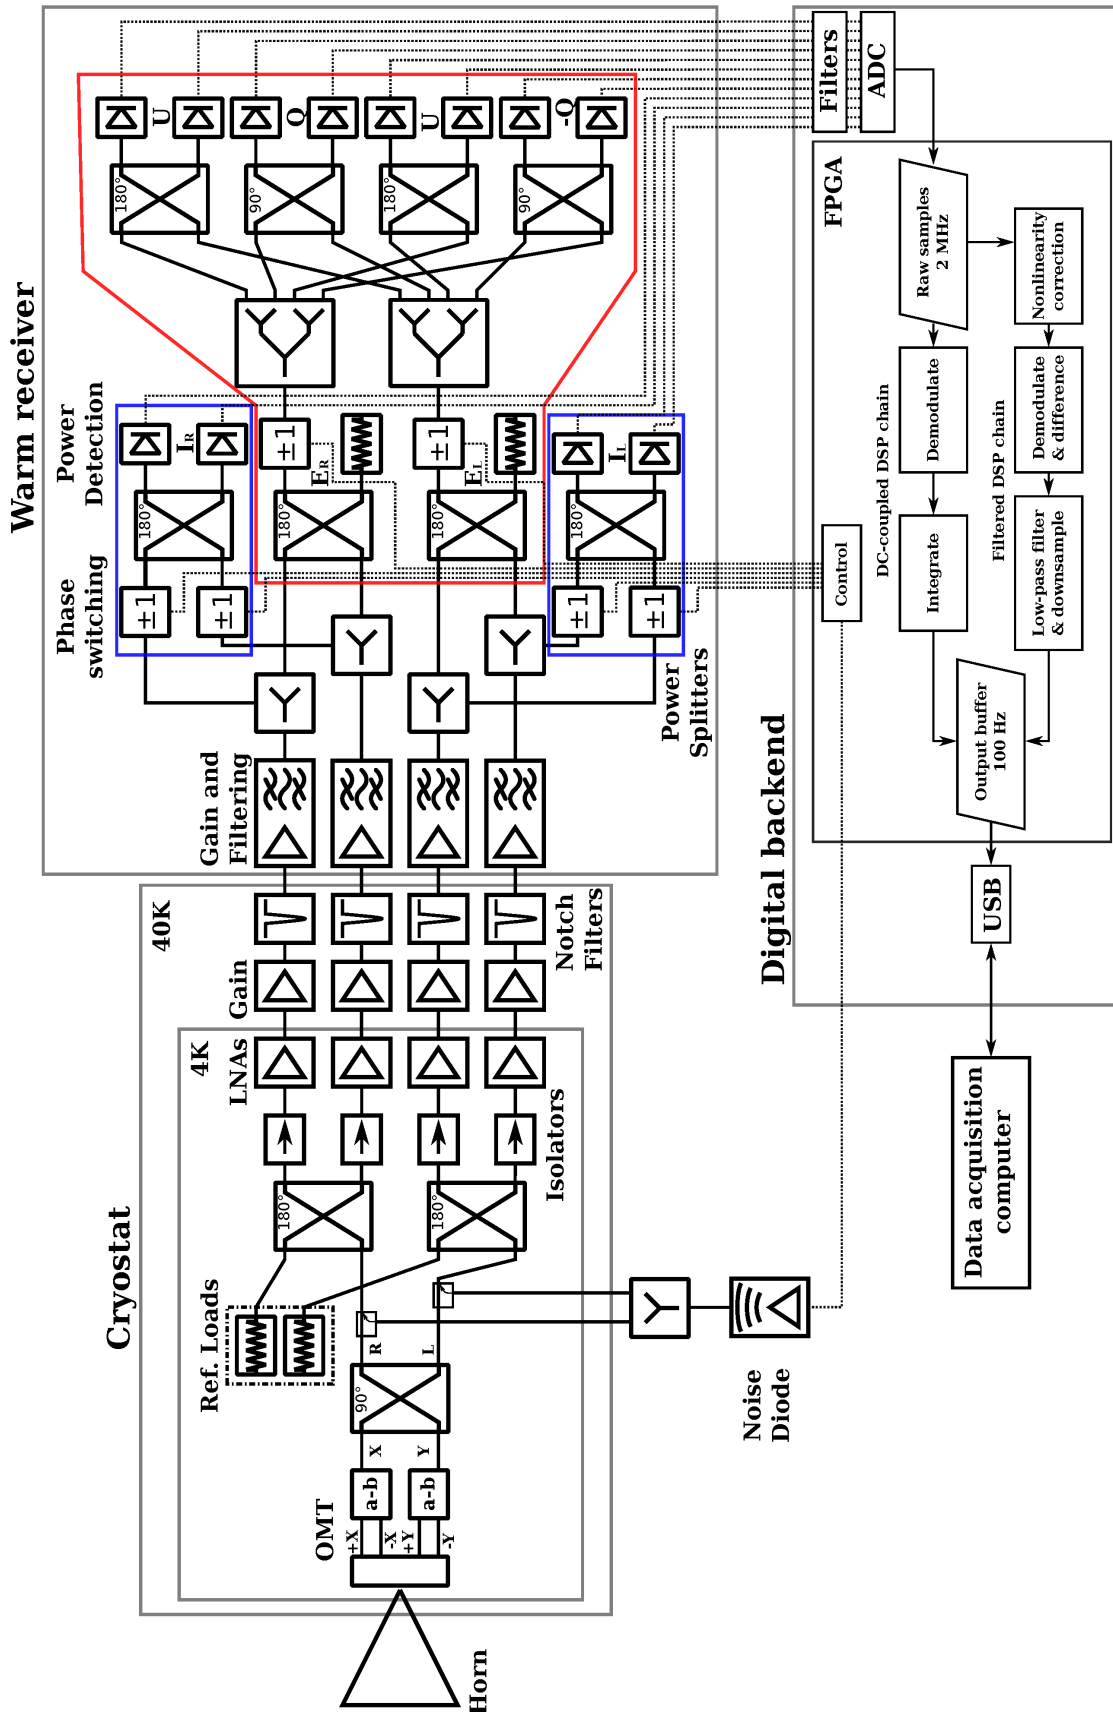
2.1 Radiometer design

The continuous-comparison radiometer architecture of Seiffert et al. (2002) reduces $1/f$ noise from receiver gain fluctuations, but without the loss of effective integration time of Dicke switch radiometers (Dicke 1946) and does so irrespective of the actual knee frequency (the frequency at which the $1/f$ noise and the thermal noise are equal). Considering only the right circularly polarized (RCP) channel, the sky signal voltage E_R and the reference signal E_{ref} are combined using a $180°$ hybrid to form sum and difference voltages $(E_R \pm E_{\mathrm{ref}})/\sqrt{2}$ (where we chose units such that $E^2 = T$, the antenna temperature). These voltages are amplified using identical gain chains with power gain G. The two signals are then multiplied and integrated. A change in gain ΔG will cause the output to vary by (Rohlfs & Wilson 2004)

$$\frac{\Delta T}{T_{\mathrm{sys}}} = \frac{\Delta G}{G}\frac{T_R - T_{\mathrm{ref}}}{T_{\mathrm{sys}}}. \qquad (2)$$

i.e. if $T_R = T_{\mathrm{ref}}$, then $1/f$ gain fluctuations will have no effect on the measured output. This comes at a price however: the level of thermal noise is $\sqrt{2}$ higher in a continuous-comparison radiometer than in a basic radiometer because of an additional amplifier and the differencing of T_R with the reference signal T_{ref}.

The multiplier in the radiometer is implemented using a second $180°$ hybrid with detector diodes on each output. The signals from the diodes are sampled, integrated, and subtracted in software to obtain the multiplication product. An advantage of this method is that

---

[2] The QUIET receiver had some pixels connected in a differential-temperature assembly that measured the temperature difference between two horns. However, this introduces spatial filtering, as discussed in Section 2.1.

Figure 1. A simplified schematic diagram of the C-BASS receiver. Left and right circular polarizations (L and R respectively) are extracted from the feed-horn and OMT and combined with temperature-controlled reference loads using 180° hybrids. After gain and filtering a bank of 180° hybrids correlates the total intensity signals (blue boxes) and recovers the circular polarization voltages, which are fed in to the correlation polarimeter section (red box). The detected powers are digitized and processed by digital signal processing (DSP) chains in an FPGA-based digital readout system and then stored on disk.

we have access to the sky and load powers T_R and T_{ref} individually in software. This is important for calibration and diagnostics.

For the C-BASS instrument we implement a continuous-comparison radiometer for each of the two orthogonal circular polarizations. The reference signals are provided by two temperature-stabilized thermal loads. In terms of Stokes parameters, the powers in the left and right orthogonal circular polarizations are $(I + V)/2$ and $(I - V)/2$ respectively, so the sum of the two measured powers is proportional to Stokes I. The difference can also be used to measure Stokes V. Since the load temperatures $T_{ref\,1}$ and $T_{ref\,2}$ are stable (see Section 3.3), variations in the final measured quantity, $I - (T_{ref\,1} + T_{ref\,2})/2$, represent the true sky brightness variations.

Some continuous-comparison receivers (e.g., WMAP, Jarosik et al. 2003) use a second feed, pointed at a different part of the sky, as the reference. This has the advantage that the receiver is well balanced, i.e. the signals being correlated have very similar temperatures, as the atmosphere (not relevant to WMAP) and CMB contributions to the sky temperature will be common to both horns. However, using a second horn to provide the reference load is unsuitable for the C-BASS experiment, which aims to measure the diffuse emission on large angular scales. Any signal common to both horns, i.e. sky structure on angular scales larger than the separation of the horn beams, is lost because a continuous-comparison receiver measures the difference in horn powers. If the second feed were feeding the same telescope optics, the maximum separation of the beams would be only a few beamwidths, and all information on scales larger than a few degrees would be lost. If instead we used a second feed-horn pointed at a large angle to the telescope axis, the large beam of the reference feed-horn compared to the full telescope optics would introduce ground contamination.

The reference load can also be a thermal source at approximately the same physical temperature as the antenna temperature (see e.g., *Planck* LFI Davis et al. 2009). A matched load, or termination, at about the same physical temperature as the sky signal can provide such a reference signal. It should be kept at a constant temperature to ensure that any variation in the output is due to the sky signal varying, and not the reference load. The C-BASS reference loads are two matched, stabilized loads housed in the receiver cryostat. These are discussed in more detail in Section 3.3. The temperature of the loads can be set to give optimum suppression of gain fluctuations, and is measured by an accurately calibrated thermistor that can also be used to help establish the temperature scale of the instrument.

2.2 Polarimeter design

Correlation polarimeters, in which the orthogonal electric field modes are correlated, are the standard polarimeter architecture at radio wavelengths, e.g., Bergano et al. (2011). Defining correlation as $X * Y \equiv \langle XY^* \rangle$, we can obtain the Stokes parameters thus (Born & Wolf 1964):

$$
\begin{aligned}
I &= E_L * E_L + E_R * E_R \\
Q + iU &= 2E_R * E_L \\
V &= E_L * E_L - E_R * E_R,
\end{aligned} \tag{3}
$$

i.e., the linear polarized Stokes parameters Q and U are the real and imaginary parts of the correlation between left and right circular polarizations. Receivers that implement the correlation approach are immune to amplifier $1/f$ noise as it is uncorrelated between amplifiers. Additive noise (e.g. the thermal noise from the amplifier) and correlated noise (e.g. due to common physical temperature changes of the amplifiers) are not suppressed.

In the C-BASS polarimeter, the first stage is to recover the pure circular polarization signals from the cryostat outputs, which are linear combinations of the right and left circular polarizations with their respective load signals. As the polarization measurement is immune to $1/f$ noise, the reference load signals are unnecessary and would add excess noise to it. The combined sky and load signals are each passed through a second $180°$ hybrid, where the reference signals $E_{ref\,1}$ and $E_{ref\,2}$ are discarded. The recovered circular polarizations E_L and E_R are phase switched and correlated to obtain two measurements each of Stokes Q and U. The multiplying elements of the correlator are implemented using hybrids to combine the signals, and detector diodes to form the squares of these sums, as in the radiometer. The complex products are formed as follows. The outputs of the $180°$ hybrids are proportional to

$$ E_L + E_R, E_L - E_R $$

and so the outputs from the detector diodes are

$$ \langle E_L^2 \rangle + \langle E_R^2 \rangle + 2\langle E_L E_R \rangle, \langle E_L^2 \rangle + \langle E_R^2 \rangle - 2\langle E_L E_R \rangle. $$

Taking the difference between these outputs gives a signal proportional to the product $\langle E_L E_R \rangle$, which is the real part of the complex correlation, $Q = 2\Re\langle E_L E_R \rangle$. Similarly, using $90°$ hybrids, the differenced output of the diodes is proportional to $\langle iE_L E_R \rangle$, which gives the imaginary part, $U = -2\Im\langle E_L E_R \rangle$. Any residual component of the total power terms $\langle E_L^2 \rangle$ and $\langle E_R^2 \rangle$ that is not perfectly removed, due to imbalances in the hybrids and post-hybrid hardware, is removed by the phase switching described in Section 2.3.

The detector output voltages are routed to the digital readout system that performs further signal processing, described in Section 4.4. The final data product is read from the digital readout system by a data acquisition computer and stored to disk.

2.3 Phase Switching

In both the radiometer and the polarimeter, phase switches are used to invert the phases of the signals prior to the correlating element. Phase switching effectively swaps the outputs of each hybrid, resulting in the sky and load signals at the output of the differencing being switched in sign. Phase switches (see Section 4.1) are placed in both arms of each pair of signals, and are driven by digital IO pins on the backend FPGA (see Section 4.4). One phase switch in each pair is switched by a $500\,\mathrm{Hz}$ square wave signal, while the second is driven by another $500\,\mathrm{Hz}$ square wave offset from the first by $90°$ of phase. The switches are cycled between all four combinations of states at an effective phase switch rate of $1\,\mathrm{kHz}$, with each sample of the integrated output signal containing data from equal times in each state. This ensures that any gain differences and offsets between the physical channels are differenced out in the integrated data. The phase switch frequency was selected to be higher than the $60\,\mathrm{Hz}$ mains signal, while still being low enough to avoid significant attenuation of the high-order harmonics of the
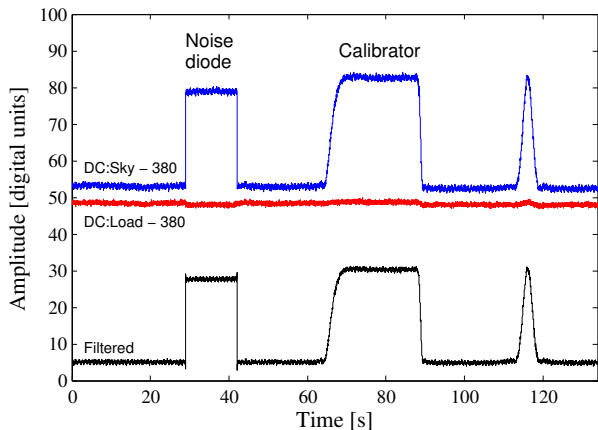
Figure 2. Total intensity data from the filtered and DC-coupled data streams described in Section 2.3. In this stretch of data the noise diode is fired, then the telescope is pointed at a calibrator source. The blue trace shows data from the phase switch state in the filtered data stream that produces the sky signal at the diode output, while the red trace shows the phase switch state that produces the load signal. These traces have been offset as indicated for clarity. The black trace is the corresponding section of the filtered data stream, which is a filtered version of the difference between the blue and red traces.

phase switch signal by the video bandwidth (VBW) of the detector diodes.

Since phase switching moves the signal of interest from the low-frequency part of the spectrum ($\lesssim 20$ Hz) to kHz frequencies, it is possible to filter out lower frequency contaminating signals, such as 60 Hz mains pickup, before demodulating and recovering the sky signal. While it is possible to do this high-pass filtering in hardware, this has the effect of removing the absolute level of the diode output, and thus losing important diagnostic information about the state of the receiver. Instead, we directly sample the DC-coupled voltages from the detector diodes, and perform two parallel reductions of the data in real-time firmware (known as DC-coupled and filtered modes). In the DC-coupled mode, the data are simply averaged in to corresponding phase switch states for each diode. This means that the sky and load signals are available independently, providing diagnostic information about the power levels, but containing some contamination, particularly from mains pickup. In filtered mode, the data are high-pass filtered, demodulated and differenced before averaging. This loses information about the individual diode signals, as only the difference (sky − load) is stored, but provides the best quality science data with 60 Hz pickup and other potential low-frequency contaminant signals removed. A short period of total intensity data from these data streams is plotted in Fig. 2, illustrating the relationship between the two data streams. Both data streams (DC-coupled and filtered) are recorded by the readout computer (Section 4.4).

2.4 Noise diode calibration

A noise diode is used to calibrate both the polarimeter and radiometer. A broadband, temperature-stabilized noise diode signal is split and coupled in to the right and left circular polarization signals in the cryostat using −30 dB directional couplers. The coupled noise power is equivalent to an antenna temperature of about 3 K and is therefore detected with a large signal-to-noise ratio of

$> 200 : 1$ in a 10 ms integration, but without changing the total system power by more than 10%. The noise diode can be switched on and off under control of the digital backend. The noise signal is coupled in to the sky signal path and should appear purely in the sky signal data channel. A detection of the noise diode signal in the load data channel is an indication of amplitude or phase imbalance between the first and second hybrids. The typical level of this leakage can be seen in the data plotted in Fig. 2. We use this as a diagnostic only, as the full phase switching regime automatically cancels out the effect of such imbalances. The amplitude of the noise diode signal is used to track gain variations in the system between astronomical calibration events.

In the polarimeter, the noise diode should generate a pure Q signal in the instrument reference frame, since it is fully correlated between the two circular polarizations. Presence of the diode signal in U indicates phase imbalance between the circular polarizations. This is monitored by regularly firing the noise diode and the effect is calibrated out of the data. The noise diode is typically fired on minute time-scales to track receiver gain fluctuations. The amplitude of the noise diode signal is calibrated against astronomical sources of known (and stable) flux density on time-scales of hours. A consequence of the noise diode injection system is that there is a small correlated noise signal injected even when the diode is switched off, due to thermal noise in the diode, and this produces a small offset in the Q channels of ≈ 30 mK that is removed in subsequent calibration.

3 CRYOGENIC RECEIVER

3.1 Cryostat design

The cryostat cools the low-noise amplifiers (LNAs) and pre-LNA components to reduce the system temperature, and to provide a cold load at comparable temperature to the sky. A 2-stage Sumitomo SRDK-408D2 Gifford-McMahon cryocooler driven by a CSA-71A air-cooled compressor is used to provide cooling. The cryostat is cylindrical, coaxial with the feed-horn. The corrugated parts of the feed-horn are bolted directly to the cryostat top plate, while the smoothly-tapered horn throat sections are machined in to the cryostat body and the internal 40 K heat shield (see Fig. 3). Gaps of 0.5 mm between the stages provide thermal isolation without compromising radio frequency (RF) performance. A Mylar window provides vacuum isolation, and a plug of Plastazote LD45 foam glued in to the 300 K smooth horn section provides infra-red blocking.

A circular copper cold plate is mounted directly to the 4 K cold head. All the 4 K components are mounted on this plate. Stainless steel coaxial cables are used to transfer the RF signals from the LNA outputs to the amplification and notch filtering (Section 3.7) on the 40 K stage, and then from the 40 K stage to the 300 K outer cryostat wall.

3.2 Orthomode transducer

The orthomode transducer (OMT), which extracts orthogonal linear polarization voltages from circular waveguide, is attached to the 4 K stage in the cryostat. It must thus be compact and easily coolable to 4 K; these conditions are not met by most commercial OMTs. We devel-
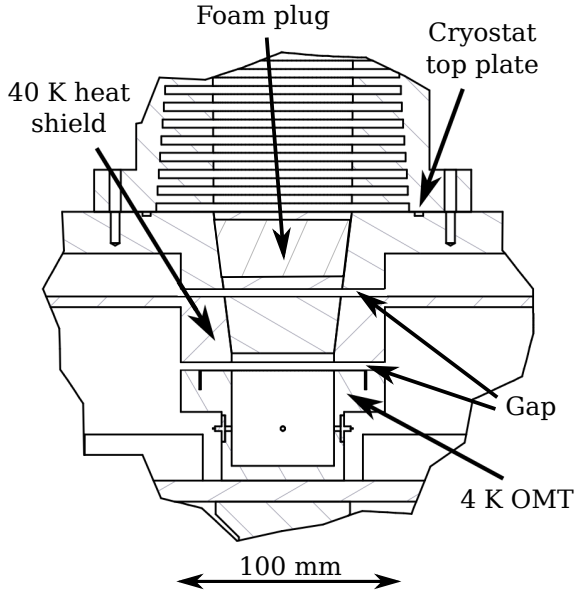
Figure 3. A cross-section of the optical feedthrough on the cryostat. The smoothly-tapered section of the horn throat is incorporated in to the top wall of the cryostat and the 40 K heat shield. A Mylar sheet (not shown) between the corrugated horn section and the cryostat body provides a vacuum window. A 0.5 mm vacuum gap (exaggerated in the drawing) separates the 300 K stage and the 40 K thermal shield. A similar break separates the 40 K heat shield and the 4 K OMT.

oped an OMT based on a design by Bock (1999), comprising four probes at right angles in a cylindrical waveguide (Grimes et al. 2007). The OMT body is machined from aluminium and contains four rectangular probes, chemically etched from a 0.1 mm thick copper sheet, placed orthogonally in a single cross-sectional plane of a circular waveguide. The probes are supported only by the pins of the SMA connectors that protrude through the waveguide wall, as shown in Fig. 4. There is a narrow (0.5 mm) break in the waveguide between the top of the OMT and the start of the horn, which is incorporated in to the 40 K heat shield. Simulations conducted in Ansoft's HFSS software[3] suggest that the effect of the break on the cross-polarization of the OMT is negligible, increasing it from −76 dB to −75 dB at 5 GHz.

The bases of the probes are tapered at an angle of 45° to the waveguide walls to reduce the capacitive coupling to the waveguide wall and are soldered to the pins of the SMA jacks, which are grounded to the body of the waveguide and fed through the wall in front of a fixed backshort. The OMT has two outputs for each linear polarization, which must be combined externally in anti-phase in the linear-to-circular converter (see Section 3.5) to give the full output signal. The signals from each pair of probes are transferred by equal-length, semi-rigid cables to the linear-to-circular converter to be combined to produce the orthogonal circular polarization signals.

Figure 4. A top view of the C-BASS OMT, showing four rectangular probes suspended orthogonally in the waveguide.
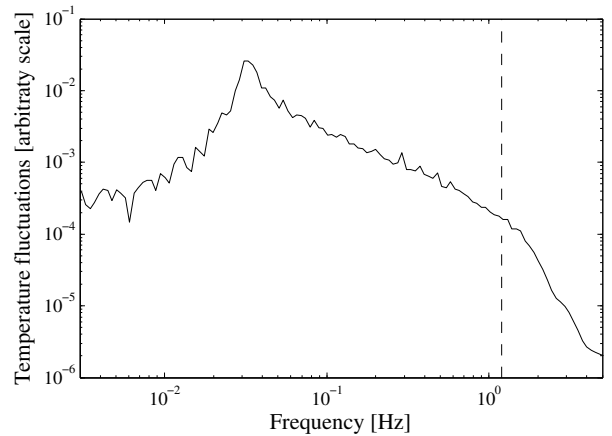

Figure 5. Power spectrum of the cold load temperature sensor data. Note the low level of mechanical refrigeration cycle signal at 1.2 Hz (indicated by the dashed line). The fluctuations have an rms of ≈ 0.4 mK.

3.3 Cold loads and temperature stabilization

The reference loads are shunt resistors attached to a copper bobbin that is mounted on the top of the cold plate. A heater coil is wound round the bobbin and a Cernox CX-1050-AA temperature sensor is embedded in it. The bobbin is mechanically bolted to the 4 K cold plate, separated from it by Teflon washers to ensure a weak thermal link. The thermal link is strong enough to cool the bobbin (giving the heater something to work against), but weak enough that the temperature of the bobbin does not greatly affect the cold plate temperature. The bobbin can be heated to 80 K without raising the cold plate temperature by more than 1 K. The reference loads are connected to the 180° hybrids through stainless steel coaxial cables, which also minimize the thermal coupling between the bobbin and the cold plate.

A Cryocon 32B temperature controller, operating as a proportional-integral-derivative (PID) control loop, is used to keep the bobbin temperature stable. A major potential contaminant of the receiver data stream is the
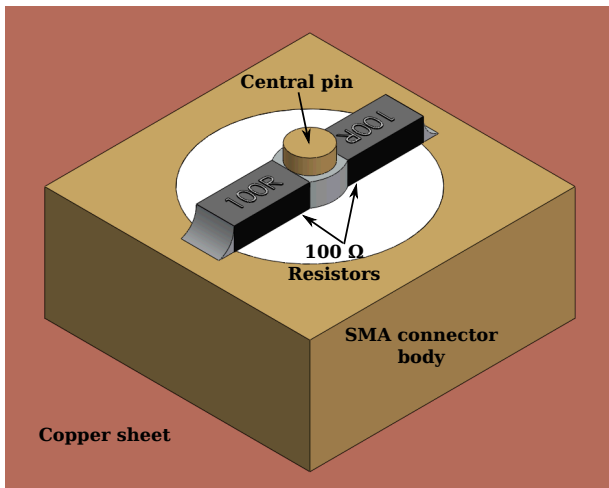
Figure 6. The C-BASS cold reference loads, rear view. An SMA straight PCB receptacle is soldered in to a copper sheet that is attached to a copper bobbin, whose temperature is controlled. Two shunt $100\,\Omega$ resistors are soldered between the central pin and outer conductor of the coaxial line. This design ensures that the resistive element of the termination is at the same physical temperature as the copper bobbin.

thermal variation due to the refrigeration cycle, which varies with a fundamental frequency of $1.2\,\mathrm{Hz}$. The measured stability of the cold load temperature sensor data is shown in Fig. 5. The refrigeration cycle is not visible in the temperature power spectrum, and the temperature is stable with an rms value of $\approx 0.4\,\mathrm{mK}$.

Initially we used commercial $50\,\Omega$ SMA RF terminations, soldered in to pockets in the copper bobbin, as the reference loads. However, the mechanical construction of the terminations meant that the shunt resistors embedded therein were weakly thermally linked to the outer body of the termination, resulting in an uncertain load temperature. We developed a second design based on SMA bulkhead connectors, in which the thermal link is stronger ensuring that the resistors are at the same temperature as the copper bobbin. In this new design two surface-mount $100\,\Omega$ resistors are soldered between the centre pin and body of the connector. This is then mounted on a small copper plate, which is in turn mounted to the top of the bobbin. A 3D CAD model of this is shown in Fig. 6. This design ensures that the resistors (which are poor thermal conductors) are well thermally connected to the bobbin. S-parameter tests with a vector network analyser confirmed that the resistors gave a good match to $50\,\Omega$ coax at cryogenic temperatures.

3.4 Hybrids

The receiver contains both $180°$ hybrids (that combine two input signals with 0 and $180°$ phase shifts) and $90°$ hybrids (that combine two input signals with $90°$ phase shifts). Both $90°$ and $180°$ hybrids are used before the LNAs in the signal chain; their loss is thus of critical importance. No commercial products had sufficiently low loss and good performance over the C-BASS band, so we designed our own hybrids.

3.4.1 180° Hybrids

Achieving a broadband $180°$ phase shift while keeping the loss low is a challenging engineering problem. The C-
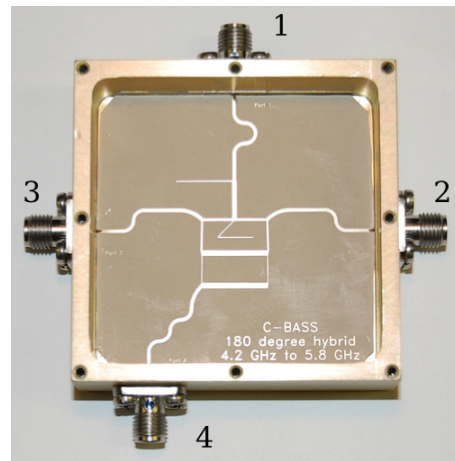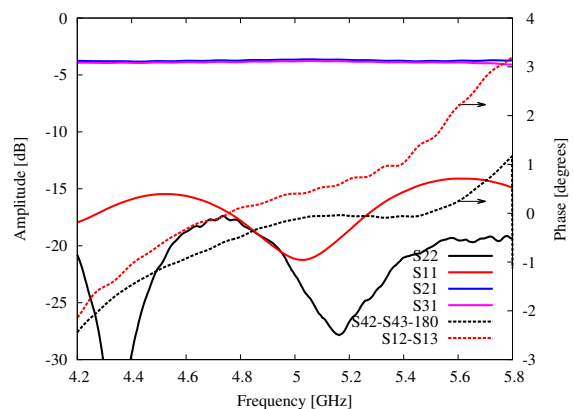
Figure 7. (a) A picture of a C-BASS $180°$ hybrid. Ports 2 and 3 are the input ports. Port 1 produces the sum of the voltages at ports 2 and 3, and port 4 produces the difference. The box housing the substrate is $50\,\mathrm{mm}$ square. (b) The measured response of the $180°$ hybrid, showing the return loss (S11 and S22), the transmission to the sum port (S21 and S31), and the phase of the sum ($\angle S12 - \angle S13$) and difference ($\angle S42 - \angle S43 - 180°$) operations. The phase shifts are within $1°\!.5$ of ideal over the full C-BASS band of 4.5 to 5.5 GHz. Amplitudes are plotted against the left-hand ordinate and phases (indicated by the arrows) against the right-hand ordinate.

BASS $180°$ hybrids are based on the 'coat hanger' model (Knochel & Mayer 1990), which is an extension of the rat-race hybrid (Pozar 2005) (the name 'coat hanger' comes from the characteristic shape of the copper traces on the substrate). The hybrid is implemented on low-loss Rogers Duroid 6010 substrate. A picture of one of the C-BASS $180°$ hybrids, and a plot of the measured response, is shown in Fig. 7.

The performance of the hybrid shown in Fig. 7(b) is close to ideal: the transmission from the inputs (ports 2 and 3) to the sum port are equal to within $0.1\,\mathrm{dB}$ over the C-BASS band. The return loss of input port 2 (S22) is typical of all the ports, and is better than $-15\,\mathrm{dB}$. The phase performance is particularly noteworthy, coming within $1°\!.5$ of ideal operation.

3.4.2 90° Hybrids

The receiver uses $90°$ hybrids in the polarimetry section to combine the circular polarizations with a $90°$ phase shift, and as circularizers in the linear-to-circular converter de-
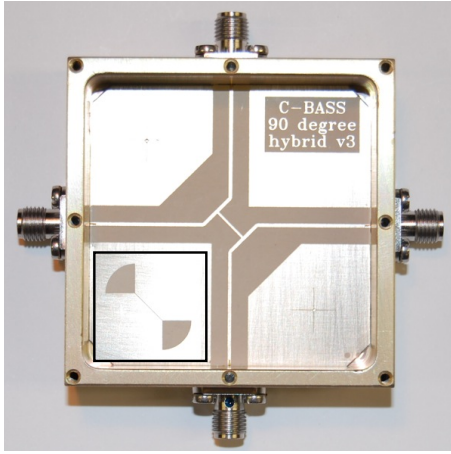
(a)
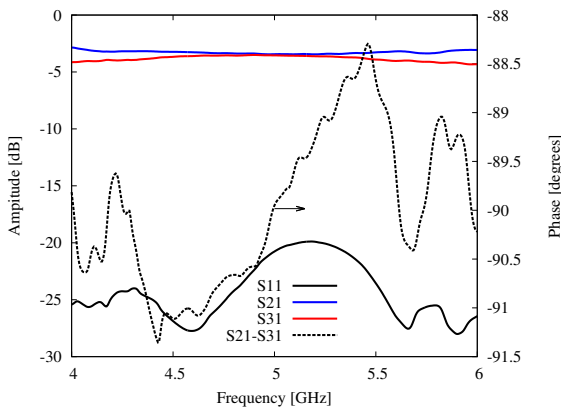
(b)

Figure 8. (a) Picture of one of the microslot 90° hybrids with the ground plane slot line shown as an inset. The slot lies directly underneath the central connecting microstrip line. The box housing the substrate is 50 mm square. (b) Measured response of a typical microslot 90° hybrid. The phase (indicated by the arrow) is plotted against the right-hand ordinate.

scribed in Section 3.5, producing orthogonal circular polarizations from the orthogonal linear polarizations produced by the OMT. Different designs were implemented in each case.

The 90° hybrids in the C-BASS polarimeter are based on the microslot coupler design (de Ronde 1970; Hoffmann & Siegl 1982a,b). They have nearly ideal phase response – within 1° of a 90° phase shift over the C-BASS band – and a return loss better than −20 dB, as shown in Fig. 8. The microslot in the bottom ground plane of the circuit lies directly underneath the central coupling line. This requires there to be a machined void in the mounting box beneath the microslot. This design is too large to fit in the space available for the linear-to-circular converter, so a different hybrid design, a microstrip branch-line coupler, is used instead.

3.5 Linear-to-circular converter

The out-of-phase signals from opposite pairs of probes in the OMT need to be combined with a 180° phase shift to obtain orthogonal linear polarizations. The linear polarizations are then passed through a 90° hybrid to produce orthogonal circular polarizations. We combine these func-

tions in to a single planar circuit that we call a linear-to-circular converter (L2C).

The signals from opposite pairs of probes are 'subtracted' (combined with a 180° phase difference) with a 3-port device based on the 180° hybrid described in Section 3.4.1 that we call a 'subtracter'. The signals from opposite probes are connected to ports 2 and 3 of the 180° hybrid, and the output signal is then seen at port 4. However, since the signals at ports 2 and 3 are 180° out of phase, port 1 is at virtual ground. Port 1 can therefore be removed entirely without affecting the performance of the device at combining 180° out of phase signals. The two 'subtracters' and branch line 90° coupler are integrated on to a single substrate and boxed. This box is attached to the cold plate close to the OMT to keep the cable length between the OMT and L2C short (10 cm).

3.6 Low noise amplifiers

The northern C-BASS receiver uses Low Noise Factory[4] LNC4_8A low noise amplifiers (LNAs), which have excellent noise $(3.3 \pm 0.4\,\mathrm{K})$ and gain $(41.6 \pm 0.3\,\mathrm{dB})$ performance between 4 and 6 GHz at cryogenic temperatures. These are preceded in the signal path by Raditek cryogenic isolators; while the isolators result in a small noise temperature penalty of $< 1\,\mathrm{K}$ (typical insertion loss is $< 0.1\,\mathrm{dB}$), they improve the matching to the LNA.

Special attention was paid to the wiring of the bias supply. Improper grounding of the wiring can result in detectable pickup of the 60 Hz supply voltage in the receiver data. The cryostat body is used as the return path for the DC bias currents. The final instrument configuration shows extremely low levels of 60 Hz pickup: it is below the noise floor of a power spectrum of 20 min of data, implying an equivalent temperature of $< 30\,\mu\mathrm{K}$.

Initially, eMERLIN C-Band LNAs were used in the cryostat. These amplifiers had typical noise temperatures of $\approx 12\,\mathrm{K}$ at 5 GHz with a gain of $\approx 30\,\mathrm{dB}$. The effect of pre-LNA losses and loading on the system by the absorbing baffles (see Holler et al. 2012 and S.J.C. Muchovej et al. in prep), along with LNA stability issues, required us to replace them with the Low Noise Factory amplifiers described here in order to achieve a satisfactory system temperature.

3.7 Notch filters

Man-made radio frequency interference (RFI) has proven to be a major challenge for the northern survey. Typical sources of RFI include aircraft radar/transponders, geostationary C-Band broadcast satellites, and fixed microwave point-to-point links (see Figure 9a for a typical unfiltered spectrum). Aircraft interference is sporadic enough that we can flag it in the data reduction pipeline (typically a few percent of data are flagged with aggressive filtering). However, the permanent nature of satellite and terrestrial RFI requires a different approach.

The C-Band satellite frequency allocation is from 3.7 to 4.2 GHz, out of the C-BASS band. However, in spite of being attenuated by 70 dB by the C-BASS bandpass filters (Section 4.3), the out-of-band satellite broadcasts were sufficiently strong that we were required to cascade

---

[4] `http://www.lownoisefactory.com/`

(a)

(b)
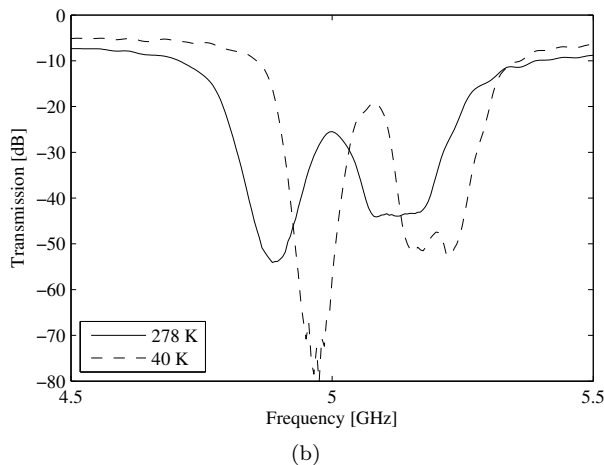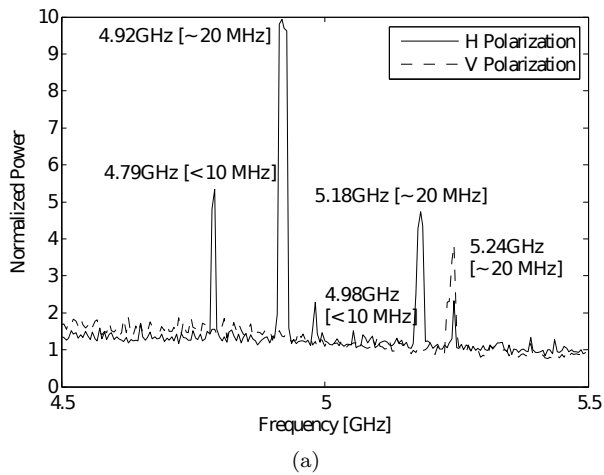
Figure 9. (a) An unfiltered measurement of the terrestrial RFI environment as measured by a spectrum analyser attached to a C-band horn. Not shown are the strong satellite broadcast signals at $\approx 4\,\text{GHz}$. The spectrum is dominated by microwave point-to-point transmitters located to the north and south of OVRO. (b) The effect of cooling on the notch filter transmission; cooling the filter to $40\,\text{K}$ increases the quality factor ($Q = \nu_c/\Delta\nu$) by reducing the resistive losses in the copper, and changes the notch centre frequencies. The depth of the notches in the $40\,\text{K}$ measurement are limited by the noise floor of the measurement equipment. The higher Q-factor is seen in the increased slope of the notch edges.

two bandpass filters (BPFs) in each signal chain to remove them from our data.

There are several terrestrial fixed links in the vicinity of the Owens Valley Radio Observatory that are bright enough to interfere with our observations. Since the northern receiver does not have a spectrometer backend, we use cryogenic notch filters tailored to the measured RFI spectrum to deal with in-band interference. This comes at the price of reduced receiver bandwidth (see Table 1), and complicated phase structure in the band (see Section 5.1).

Ideally the notch filters should have a relatively high quality factor $Q = \nu_c/\Delta\nu$, where ν_c is the centre frequency and $\Delta\nu$ the half-power bandwidth of the notch. Each RFI band is approximately $25\,\text{MHz}$ wide (see Fig. 9a), which at $5\,\text{GHz}$ requires a Q-factor of 200. A quality factor this high requires superconducting filters. However, since each of the major RFI bands (i.e. $4.92\,\text{GHz}$ and $5.18\,\text{GHz}$) seen at OVRO has a second proximate RFI peak, a somewhat broader stop-band can be used. We de-
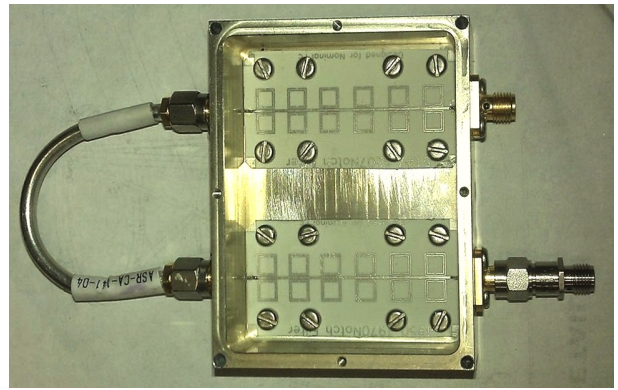
Figure 10. The C-BASS RFI notch filter assembly, consisting of two cascaded split-ring resonators housed in a single box. An isolator is placed on the input to improve the matching.

signed two notch filters, each with a stop-band of $80\,\text{MHz}$, to target these RFI bands. The notch filters are split-ring resonators, which use the electromagnetic coupling between resonant structures placed parallel to a microstrip transmission line (Garcia-Garcia et al. 2005; Futatsumori et al. 2008), and are fabricated using copper on low-loss substrate (Copley 2013). A picture of the filter assembly is shown in Fig. 10; both notch filters are housed in a single box and an isolator is included on the input to improve the matching.

The quality factor of the resonator is largely determined by the resistivity of the conductor, so cooling the filters results in a narrowing of the filter response and a small shift in the resonant frequency, as shown in Fig. 9b. We placed the filters on the $40\,\text{K}$ stage of the cryostat, which gave a good match to the desired stop-band.

4 POLARIMETER AND RADIOMETER

The warm receiver components, along with the calibration noise diode, are housed in a temperature controlled box beneath the cryostat and above the dish surface. The gain and filter chains are shown as a simplified block diagram in Fig. 1. They consist of isolators, slope compensators, a bandpass filter, and multiple amplifiers. Below we describe the design of the major receiver components; further details are in King (2009).

4.1 180° phase switches

The phase switches were constructed using two broadband surface-mount HMC547LP3 Single Pole Double Throw (SPDT) switches to select between 0° and 180° signal paths with a rise time of $3\,\text{ns}$. The phase switch is powered by a $-5\,\text{V}$ DC supply and can be switched by CMOS or TTL logic signals.

We obtain a 180° phase shift through geometric means by using a microstrip to slotline transition. The electric field in the microstrip line is coupled to a slotline in the ground plane. The signal is then coupled back in to a microstrip line at the other end of the slotline. Both microstrip lines and the slotline are terminated with radial stubs. If the second microstrip line exits on the opposite side of the slotline as compared to the input microstrip line then the electric field is effectively flipped, and a 180° phase shift is achieved. A drawing of the microstrip to

(a)

(b)
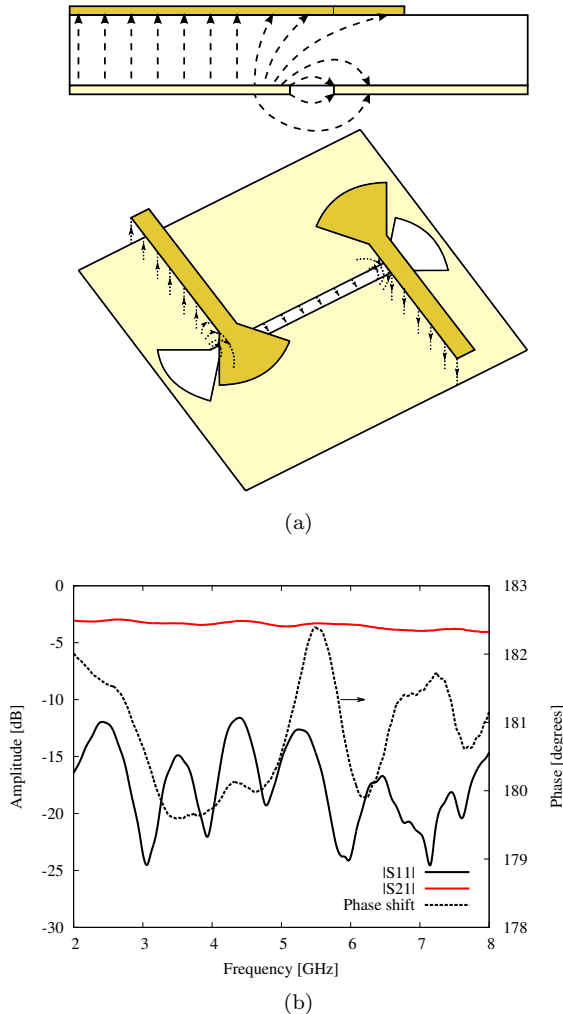
Figure 11. (a) Drawing of the transition used to achieve a 0° or 180° phase shift. The upper trace is coloured gold, the ground plane is coloured yellow, and the electric field vectors are indicated by dashed lines. (b) The response of the phase switch. The induced phase shift is within a few degrees of 180° from 2 to 8 GHz. The phase (indicated by the arrow) is plotted against the right-hand ordinate.

slotline structure, and the measured response of the phase switch, is shown in Fig. 11.

4.2   Warm amplifiers

The warm amplifiers, which provide all of the gain after the LNAs, are built around the Hittite HMC462LP5 broadband distributed amplifier monolithic microwave integrated circuit (MMIC). Each amplifier block uses two stages of this self-biased, surface-mount low noise amplifier – based on GaAs pseudomorphic high-electron-mobility transistor (pHEMT) technology – to provide ≈ 30 dB of gain from 2 to 20 GHz with a noise temperature of ≈ 200 K. The MMIC packages are interconnected using grounded coplanar waveguide rather than microstrip, as this allows for better broadband match between the devices and the SMA connectors, and a better grounding environment for the MMIC packages.

4.3   Bandpass filters

The bandpass filters consist of two cascaded filters – a broadside edge-coupled bandpass filter (BPF) followed

by a stepped impedance lowpass filter (LPF) to remove the higher harmonic images of the BPF response. The filters were fabricated on 0.254 mm thick RO4350 substrate ($\epsilon_r = 3.66$). The 3 dB bandpass of the filter is 4.5 to 5.5 GHz. Each filter provides ≈ 70 dB of rejection in the stop-band. This is not sufficient to reject the signals from the brightest C-Band direct-broadcast satellites, so two complete filter units are cascaded in each RF signal path.

4.4   Detectors and digital backend

The schematic diagram in Fig. 1 shows the layout of the detection, digitization, and digital signal processing scheme. We use zero-bias Schottky detector diodes, with a video bandwidth of around 800 kHz, to measure the RF power. The detector diode voltages are then filtered and digitized, as described in Section 4.4.1. A digital signal processing (DSP) chain (Section 4.4.2) produces the final data stream from the raw ADC values, which are transmitted over a USB link to the data acquisition computer.

The filter card and ADC/FPGA-based digital card are housed in the digital backend box behind the telescope primary mirror. The cards are connected by a backplane, powered by linear power supplies, and housed in shielded enclosures.

4.4.1   Filtering and digitization

On the filter board the voltages are DC-coupled and low-pass filtered at 1 MHz, the Nyquist frequency of the ADCs, by active voltage-control voltage-source (VCVS) anti-aliasing filters. The filtered voltages are amplified by a factor of ≈ 6.5 in order to avoid digitization noise and maximize the dynamic range of the digitization.

The digital card is designed around 16 ADCs capable of sampling at 2.77 MHz and two Xilinx Spartan 3 FPGAs. The card was originally designed for the Linear Collider And Survey (LiCAS) particle physics experiment (Reichold et al. 2006). For the C-BASS application, only 12 of the ADCs are used and sampling is performed at 2 MHz with 14 bit resolution. One of the FPGAs was reprogrammed for C-BASS specific digital signal processing (DSP) and control functions, while the other is an interface to on-board memory that is used for diagnostic purposes.

4.4.2   Digital processing

The DSP chain demodulates and integrates the twelve 2 MHz data streams to 10 ms integrations. The raw data streams contain 60 Hz mains pickup. In the DC-coupled mode used for diagnostic purposes that preserves the full channel-by-channel information (see Section 2.3), this signal and its harmonics are aliased to frequencies in the range of the science data (0.01 – 10 Hz), which would put undesirable contamination in the final images.

For the science data stream, the signals are first corrected for non-linearity of the detector diodes, using look-up tables determined from laboratory measurements of the individual diodes. The signals are then differenced to produce the (sky − load) signals, which also reduces the number of data streams from 12 to 6. The 1 kHz phase switch modulation is then demodulated, which reduces the science signal to baseband while mixing the mains

and its harmonics to higher frequencies. A chain of decimating low-pass filters follows, reducing the data rate to 100 Hz and applying a 50 Hz rectangular low-pass filter to the data.

All these functions are carried out in the FPGA, which also produces the phase switch drive signals and controls the noise diode. A USB interface transmits the integrated data to the telescope control computer and also allows control of the digital backend functions, such as firing the noise diode. A detailed description of the digital backend is given in Stevenson (2013).

5 INSTRUMENT PERFORMANCE

We have discussed the performance of individual components of the receiver throughout the text. In this section we discuss the performance of the receiver as a whole. This discussion is split in to three sections: the variation of the receiver response with frequency (i.e. the passband) in Section 5.1, the stability of the receiver in Section 5.2, and the sensitivity of the receiver in Section 5.3.

5.1 Passband

The signal measured in each I, Q, U polarization output is the integration of the instrument response across the passband. In the DC-coupled diagnostic data stream we measure the antenna temperature and reference load temperature as two separate data streams, which we can write as (in this case, for the left circular polarization (LCP) channel):

$$r_L = \int G_L(\nu) \left[(1 + \alpha_L(\nu)) T_{\rm L} + (1 - \alpha_L(\nu)) T_{\rm ref} + T_N \right] d\nu \tag{4}$$

$$r_{\rm ref} = \int G_L(\nu) \left[(1 - \alpha_L(\nu)) T_{\rm L} + (1 + \alpha_L(\nu)) T_{\rm ref} + T_N \right] d\nu. \tag{5}$$

Here r_L and $r_{\rm ref}$ denote the recorded data values, $G_L(\nu)$ is the frequency-dependent gain of the channel, and $\alpha_L(\nu)$ (which we call the imbalance parameter) parametrizes the separation of the sky signal T_L and the reference load signal $T_{\rm ref}$ by the continuous comparison radiometer architecture. T_N is the noise temperature of the receiver. The digital backend also records the filtered mode differenced data stream $r_L - r_{\rm ref} \propto T_L - T_{\rm ref}$. Ideally, $\alpha = 1$, i.e. the sky and load are perfectly separated by the receiver.

We can express the Q and U channels of the receiver in terms of the input linear polarization vector $(Q_{\rm in}, U_{\rm in})$, in the reference frame of the OMT, as:

$$r_Q = \int G_P(\nu) \left[\cos \theta(\nu) Q_{\rm in} - \sin \theta(\nu) U_{\rm in} \right] d\nu \tag{6}$$

$$r_U = \int G_P(\nu) \left[\sin \theta(\nu) Q_{\rm in} + \cos \theta(\nu) U_{\rm in} \right] d\nu, \tag{7}$$

where $G_P(\nu)$ and $\theta(\nu)$ are the frequency-dependent gain and polarization angle (electric vector position angle, EVPA) rotation of the instrument respectively. Ideally, for no loss in sensitivity, the instrument would have a flat passband ($G_P(\nu) = $ constant in the passband) and a constant EVPA rotation angle ($\theta = $ constant).

We measure the passband of the northern C-BASS instrument by injecting a sinusoidal voltage of known power with equal amplitude and phase in to the noise diode injection ports. It appears as a purely linearly polarized
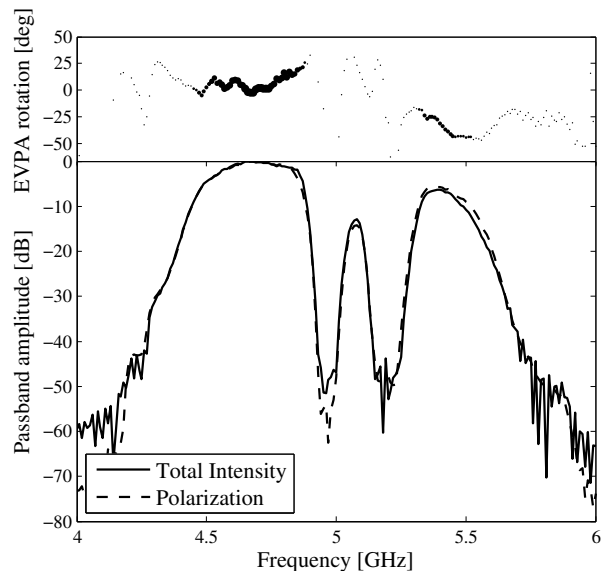
Figure 12. (bottom) The passband amplitude of the total intensity (solid) and linear polarization (dashed) channels of the C-BASS instrument. The notches from the RFI filters are clearly seen in the centre of the band. (top) The EVPA rotation angle θ of the instrument (Equations 6 and 7). The size of each frequency point is proportional to the amplitude of the instrument response (lower plot) at that frequency.

Table 1. The centre frequency and noise-equivalent bandwidth of the total intensity and linear polarization channels, assuming a flat spectrum source in power or temperature.

	Center freq. [GHz]	Bandwidth [GHz]
Total Intensity	4.764	0.489
Polarization	4.783	0.499

signal in the instrument reference frame. The frequency of this signal is swept across the C-BASS band. At each frequency point we turn the voltage signal generator on and off, and record the response of the digital backend in each state. By taking the difference of the receiver output in the on and off states we can remove the baseline signal due to the system noise from the data and directly measure the frequency-dependent response of the receiver.

The effective centre frequency and noise-equivalent bandwidth of the receiver are calculated according to the equations:

$$\nu_c = \frac{\sum_i \nu_i G_i}{\sum_i G_i}$$

$$\mathrm{BW} = \Delta \nu \frac{\left(\sum_i G_i \right)^2}{\sum_i G_i^2}. \tag{8}$$

Here ν_i, G_i are the frequency and gain of measurement i, and $\Delta \nu$ is the width of each frequency bin (the gain is assumed to be constant in each bin). This definition assumes a flat spectrum source in power (or temperature). The parameters thus calculated are shown in Table 1. Both the total intensity and linear polarization channels have very similar bandwidths and centre frequencies. The bandwidth of the receiver has been reduced by a factor of two by the notch filters.

The passband of the instrument is shown in Fig. 12, where we plot the mean of the two total intensity channels and the linear polarization intensity in the lower panel. In
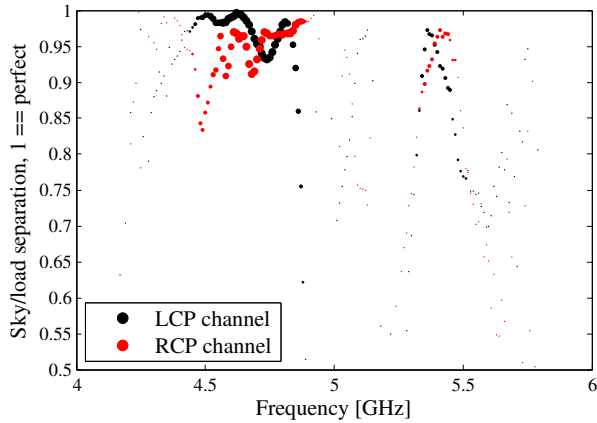
Figure 13. The separation of sky and load by the continuous comparison radiometer in the total intensity LCP (black) and RCP (red) channels, as parameterized by α in Equations 4 and 5. A value of $\alpha = 1$ indicates perfect separation, $\alpha = 0$ indicates no separation at all. The size of each frequency point is proportional to the amplitude of the instrument response at that frequency.

the upper panel we plot the EVPA rotation angle θ from Equations 6 and 7. The intensity of each point has been weighted by the amplitude of the instrument response at that frequency.

We know from measuring the passband without the notch filters in place that the large-scale variation in θ across the band – the difference between the lower part of the band and the upper part – is due to the phase structure in the notch filter transmission curves. This non-constant value of θ across the band will reduce our sensitivity. We can estimate the effect of the variation in θ by calculating the ideal and actual response of the instrument to a purely linearly polarized source. We calculate an 'ideal' instrument response by assuming $\theta(\nu) = 0$ in Equations 6 and 7. For a source with a flat spectrum, i.e. $Q_{\rm in} = Q_0, U_{\rm in} = 0$ the amplitude for the real instrument is 3.4% lower than for the ideal instrument. For a more realistic source spectrum of $Q_{\rm in} = Q_0 \nu^{-3}, U_{\rm in} = 0$ the real instrument response is 2.5% lower than the ideal instrument. This small reduction in polarization sensitivity is an acceptable price to pay for removing the RFI signature from the data.

The sensitivity of the total intensity channel is affected by the ability of the continuous comparison radiometer to properly separate the sky and load signals at the input to the power detection stage in Fig. 1. If $\alpha = 0$, i.e. the sky and load signals are perfectly mixed, then $r_L - r_{\rm ref} = 0$ regardless of the sky temperature T_L: the signal is lost entirely. The value of α for both continuous comparison radiometers is shown in Fig. 13; the intensity of each frequency measurement has been weighted by the instrument gain. The gain-weighted mean value of α is 0.96 for the LCP channel, and 0.95 for the RCP channel.

5.2 Stability

Fig. 14 shows the power spectrum of data taken by the northern C-BASS instrument while observing the north celestial pole (NCP) region where the astronomical signal remains constant due to our circularly symmetric beam (Holler et al. 2012). All the data plotted here are unprocessed and uncalibrated. This is the only measurement presented in this paper based on astronomical data. The
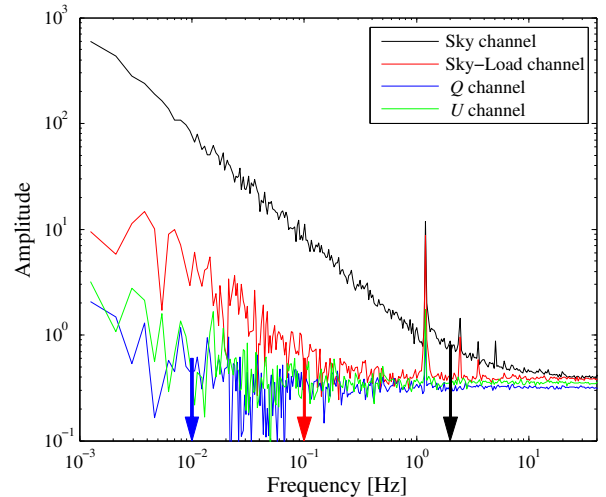
Figure 14. Power spectra of data from the northern C-BASS receiver while observing the NCP for a continuous 20 min time span. The data have had a mean subtracted but are otherwise uncalibrated. The strong 1.2 Hz refrigeration cycle signal and its harmonics are removed later by the data reduction pipeline. The black trace shows the sky data channel. The red trace shows the total intensity channel (sky−load). The knee frequencies are indicated by vertical arrows. The reduction of $1/f$ noise by a factor of ≈ 20 is testament to the power of the continuous-comparison radiometer architecture. The green and blue traces show the Stokes Q and U channels, with knee frequencies below 10 mHz.

black trace shows the sky data channel from the DC-coupled data stream, which is analogous to what a conventional radiometer would measure. The red trace shows the total intensity ($T_L - T_{\rm ref}$) data channel from the filtered mode data stream. As expected, the level of $1/f$ noise in the total intensity channel is reduced compared to the sky channel. It is a factor of ≈ 20 lower, and the knee frequency has moved from ≈ 4 Hz to ≈ 0.2 Hz. The power spectrum at 0.2 Hz is dominated by fluctuations in atmospheric emission, which we have measured to have a typical knee frequency of 0.1 Hz.

The green and blue traces show the Stokes Q and U channels, both with knee frequencies of below 10 mHz. This is sufficient given the typical scanning time of ≈ 90 s (11 mHz), and given that we remove large-scale fluctuations using destriping map-making (Sutton et al. 2009, 2010). Fig. 14 demonstrates that both the continuous-comparison radiometer and the pseudo-correlation polarimeter sections of the C-BASS receiver remove the effect of receiver gain fluctuations on the output signal.

5.3 Sensitivity

The receiver sensitivity is determined by the system temperature and the passband. The radiometer design is a compromise between thermal white noise level and low-frequency $1/f$ noise, as discussed in Section 2.1; the continuous-comparison design reduces the $1/f$ noise, but requires there to be lossy components in front of the LNAs, which increase the white noise level. We developed a receiver noise model based on the Friis noise equation (Friis 1944) to predict the noise temperature of the receiver and verify its performance. The model uses measured or estimated losses and physical temperatures for every component in the receiver, from the feed-horn through to the second-stage (warm) amplifier. The noise

contribution of subsequent components is negligible. The model will be described in detail in S.J.C. Muchovej et al. (in prep.).

The noise temperature was determined by measuring the ratio of powers at the output of the cryostat when placing hot and cold loads in the beam of the feed-horn, commonly known as a Y-factor $Y = P_{\rm hot}/P_{\rm cold}$ (Kraus 1986), where $P_{\rm hot}$ and $P_{\rm cold}$ are the measured powers given a hot and cold load respectively. While the noise model is sufficiently complicated to warrant publication elsewhere, it is constructive to consider a simplified case, as follows. For a simple receiver model in which the total noise temperature can be modelled as the sum of a receiver temperature $T_{\rm rx}$ and the hot or cold load temperatures $T_{\rm hot}$ or $T_{\rm cold}$, the receiver temperature can then be calculated from

$$T_{\rm rx} = \frac{T_{\rm hot} - Y T_{\rm cold}}{Y - 1}. \qquad (9)$$

For the continuous-comparison architecture the situation is complicated by the fact that each RF output from the receiver is the sum of voltages from a sky channel and the internal load. A simplified model considers the input to the receiver as being the mean of the load temperature and the temperature seen by the feed, resulting in the modified relation

$$T_{\rm rx} = \frac{(T_{\rm hot} + T_{\rm load})/2 - Y(T_{\rm cold} + T_{\rm load})/2}{Y - 1}. \qquad (10)$$

This equation makes the assumption that the overall gain from the load through to the receiver output is the same as that from the feed-horn aperture to the output. This is not strictly true, and for detailed comparison with the noise measurements we use the full receiver model. The purpose of this equation is to present the reader with an idea of how the various parameters affect the receiver temperature.

We used the sky as the cold load with the telescope pointed at zenith, and a microwave absorbing foam at ambient temperature as the hot load. The sky temperature at 5 GHz is a weak function of atmospheric conditions under typical conditions at Owens Valley, and on a clear dry day is typically ≈ 5 K (e.g. Smith 1982), including 2.7 K from the CMB. The precipitable water vapor at OVRO is typically 3 to 8 mm on a clear day (Philhour 2002). We made two sets of measurements, one measuring the receiver power using a power meter, with a bandpass filter to set the same bandwidth as seen by the radiometer, the other using a spectrum analyser to capture the power as a function of frequency. The former method provides a close analogue of the power detected by the broadband polarimeter/radiometer, while the latter allows us to check for any anomalous behaviour as a function of frequency.

The broadband Y-factor measurements for the four receiver output channels without the secondary optics in place produce receiver temperatures, using Equation 10 with $T_{\rm load} = 28$ K, of 22, 20, 16 and 20 K, with a mean value of 19.5 ± 2.2 K. The spectral measurements of the system temperature, calculated using Equation 10, are shown in Fig. 15. A more accurate analysis using the full receiver model yields a mean receiver temperature (not including the sky temperature and optics contribution) of 16.0 ± 1.8 K. This includes cold losses estimated from warm laboratory measurements of individual components, where we find a loss for the sky channel before the hybrid of -1.7 ± 0.5 dB (comprising the OMT, the
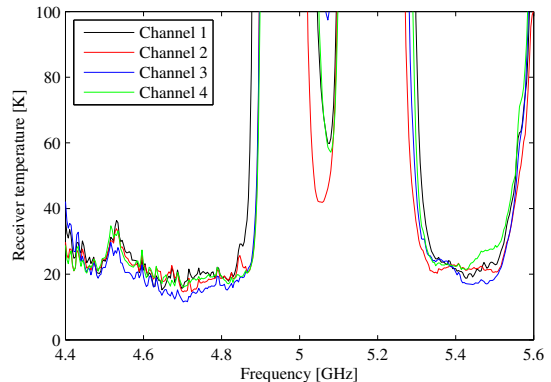
Figure 15. Measured receiver temperature of the C-BASS receiver as a function of frequency, for each of the four cryostat RF output channels. The receiver temperatures are derived from power measurements taken with a spectrum analyser with hot and cold loads, using Equation 10. The effect of the notch filters can clearly be seen. The effective receiver temperature is the gain-weighted mean of this plot.

L2C, the noise coupler, and all their interconnecting cables), a loss between the cold load and the hybrid of -1.6 ± 0.1 dB, and a cold loss common to both signal paths (including the insertion loss of the hybrid itself, the isolators, and interconnecting cables) of -1.4 ± 0.2 dB. The directly measured losses are quite uncertain (due to uncertainty in how they change with temperature), so our best estimates of the losses come from the actual noise temperature measurements.

The expected system temperature in polarization will be the receiver temperature, plus the contribution from the sky (CMB and atmosphere), the contribution from the telescope optics, and spillover. This is expected to total ≈ 30 K.

The effective system temperature in intensity is more complicated, since it has to take in to account the noise contribution from the cold load as well as the sky channel. Since each intensity output is the difference of the sky channel and the load channel, and these have independent noise signals, the noise level on one intensity output is $T_{\rm I} = \left(T_{\rm sky\ channel}^2 + T_{\rm load\ channel}^2 \right)^{1/2}$, where $T_{\rm sky\ channel} = T_{\rm sky} + T_{\rm spillover} + T_{\rm Rx\ sky}$, and $T_{\rm load\ channel} = T_{\rm load} + T_{\rm Rx\ load}$. $T_{\rm Rx\ sky}$ and $T_{\rm Rx\ load}$ are the receiver contributions to the sky and load channels respectively; $T_{\rm Rx\ sky} \neq T_{\rm Rx\ load}$ because of the different losses along the two signal paths. The noise on the final summed intensity channel is then reduced by a factor $\sqrt{2}$ as the individual intensity channels have independent noises. For a load temperature of 24 K, which minimizes the $1/f$ fluctuations in the intensity data, this gives an expected effective system temperature in intensity of 32 K.

6 CONCLUSIONS

The C-BASS project aims to produce all-sky maps of total intensity and linear polarization at moderate resolution ($0.^{\circ}73$) with a sensitivity of 0.1 mK per pixel. These will be used primarily for removal of foregrounds from measurements of the polarized CMB and will also be used, among other things, to study the Galactic magnetic field, to constrain the frequency spectrum of the so-called 'anomalous' emission and to study the Galactic cosmic ray population.

We have described the instrument for the northern

C-BASS survey which is a novel hybrid of a continuous-comparison radiometer and a correlation polarimeter. A cryogenically cooled front end allows sensitive measurements to be made. Most of the receiver components have been specifically designed to provide excellent performance in our band. A digital backend digitizes and processes the detected powers.

We limit the level of undesirable spurious signals, such as pickup from the 60 Hz supply voltage, by using phase switches to remove any low-frequency fluctuations in the detector and digitization hardware. We have shown that the continuous-comparison architecture works as designed by reducing the $1/f$ noise seen in the total intensity channel, and that the polarization channels show the expected very low $1/f$ noise with a knee frequency of 10 mHz. The combination of low $1/f$ noise and a system temperature of ≈ 30 K in both total intensity and linear polarization will allow us to reach our target sensitivity of 0.1 mK per pixel.

REFERENCES

Bergano M., Fernandes F., Cupido L., Barbosa D., Fonseca R., Ferreira I., Grossan B., Smoot G., 2011, Experimental Astronomy, 30, 23, doi:10.1007/s10686-011-9217-6

Bock D., 1999, BIMA Memo No. 74

Booth R. S., Jonas J. L., 2012, African Skies, 16, 101, ADS:2012AfrSk..16..101B

Born M., Wolf E., 1964, Principles of Optics, 2nd edn. Pergamon Press, archive.org

Buder I., 2010, Millimeter, 7741, 38, doi:10.1117/12.857882

QUIET Collaboration, 2012, arXiv:1207.5562

Copley C., 2013, PhD thesis, University of Oxford

Davis R. J. et al., 2009, J. Instrum., 12, 12002, doi:10.1088/1748-0221/4/12/T12002

Dicke R. H., 1946, Review of Scientific Instruments, 17, 268, doi:10.1063/1.1770483

Friis H. T., 1944, in Proceedings of the IRE. pp 419–422, doi:10.1109/JRPROC.1944.232049

Futatsumori S., Hikage T., Nojima T., Akasegawa A., Nakanishi T., Yamanaka K., 2008, in Proceedings of the 38th European Microwave Conference. pp 1145–1148, doi:10.1109/EUMC.2008.4751661

Garcia-Garcia J. et al., 2005, Microwave and Optical Technology Letters, 44, 376, doi:10.1002/mop.20640

Ghosh T., Banday A. J., Jaffe T., Dickinson C., Davies R., Davis R., Gorski K., 2012, Monthly Notices of the Royal Astronomical Society, 422, 3617, doi:10.1111/j.1365-2966.2012.20875.x

Gold B. et al., 2011, The Astrophysical Journal Supplement, 192, 15, doi:10.1088/0067-0049/192/2/15

Grimes P. K., King O. G., Yassin G., Jones M. E., 2007, Electronics Letters, 43, 1146, doi:10.1049/el:20071649

Hoffmann R. K., Siegl J., 1982a, Microwave Theory and Techniques, IEEE Transactions on, 30, 1205, doi:10.1109/TMTT.1982.1131222

Hoffmann R. K., Siegl J., 1982b, Microwave Theory and Techniques, IEEE Transactions on, 30, 1211, doi:10.1109/TMTT.1982.1131223

Holler C. et al., 2012, Antennas and Propagation, IEEE Transactions on, 99, 1, doi:10.1109/TAP.2012.2219843

Jarosik N. et al., 2003, Astrophysical Journal Supplement Series, 145, 413, doi:10.1086/346080

King O. G., 2009, PhD thesis, University of Oxford

King O. G. et al., 2010, in Holland W. S., Zmuidzinas J., eds, SPIE Astronomical Telescopes and Instrumentation: Observational Frontiers of Astronomy for the New Decade. SPIE, 77411I, doi:10.1117/12.858011

Knochel R., Mayer B., 1990, IEEE Transactions on Microwave Theory and Techniques Symposium Digest, pp 471–474 vol.1, doi:10.1109/MWSYM.1990.99621

Kogut A. et al., 2007, The Astrophysical Journal, 665, 355, doi:10.1086/519754

Kraus J. D., 1986, Radio Astronomy, 2nd edn. Cygnus-Quasar Books

Macellari N., Pierpaoli E., Dickinson C., Vaillancourt J. E., 2011, Monthly Notices of the Royal Astronomical Society, 418, 888, doi:10.1111/j.1365-2966.2011.19542.x

Page L. et al., 2007, Astrophysical Journal Supplement Series, 170, 335, doi:10.1086/513699

Philhour B., 2002, PhD thesis, California Institute of Technology

Pozar D., 2005, Microwave engineering, 3rd edn. John Wiley & Sons, Inc. New York

Reichold A., Dawson M., Green J., 2006, Proc. of EPAC06, pdf

Rohlfs K., Wilson T. L., 2004, Tools of Radio Astronomy, 4th edn. Springer

Seiffert M., Mennella A., Burigana C., Mandolesi N., Bersanelli M., Meinhold P., Lubin P., 2002, Astronomy and Astrophysics, 391, 1185, doi:10.1051/0004-6361:20020880

Smith E. K., 1982, Radio Science, 17, 1455, doi:10.1029/RS017i006p01455

Stevenson M. A., 2013, PhD thesis, California Institute of Technology

Sutton D., Johnson B. R., Brown M. L., Cabella P., Ferreira P. G., Smith K. M., 2009, Monthly Notices of the Royal Astronomical Society, 393, 894, doi:10.1111/j.1365-2966.2008.14195.x

Sutton D. et al., 2010, Monthly Notices of the Royal Astronomical Society, 407, 1387, doi:10.1111/j.1365-2966.2010.16954.x

Wolleben M., Landecker T. L., Reich W., Wielebinski R., 2006, Astronomy and Astrophysics, 448, 411,

doi:10.1051/0004-6361:20053851

de Ronde F. C., 1970, Microwave Symposium, G-MTT 1970 International, pp 184–189, doi:10.1109/GMTT.1970.1122803