

First Season COMAP Results: Data Processing

MARIE K. FOSS,¹ HÅVARD T. IHLE,¹ JOWITA BOROWSKA,¹ HANS KRISTIAN ERIKSEN,¹ JONAS LUNDE,¹ NILS-OLE STUTZER,¹
DUNCAN WATTS,¹ INGUNN K. WEHUS,¹ AND THE COMAP COLLABORATION

¹*Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, N-0315 Oslo, Norway*

ABSTRACT

We describe the data processing pipeline used to analyze the first season COMAP observations, converting raw detector readouts to calibrated sky maps. This pipeline implements five main steps, namely gain calibration, pointing reconstruction, filtering and data selection, noise characterization, and mapmaking. Absolute gain calibration relies on a combination of hardware and astrophysical sources, while relative gain calibration exploits real time total power variations. High efficiency filtering is achieved through common-mode spectroscopic rejection within and across receivers, resulting in nearly uncorrelated white noise within single-frequency channels. Consequently, near-optimal maps are produced by binning the filtered time stream into a pixelized map structure. Error propagation and statistical characterization are established through end-to-end simulations, and residual systematics are constrained through null-tests. After these steps, we find that 85% of the total data volume is retained for scientific analysis, leading to an average map sensitivity of $20\ \mu\text{K arcmin}$. Systematic uncertainties are limited to at most 20% of the standard deviation of the final map.

Keywords: stars: formation, galaxies: star formation, radio lines: galaxies, methods: data analysis

Contents		3.4.1. Calibration using a reference load	12
		3.4.2. Photometric Calibration	12
1. Introduction	1	3.5. Noise characterization	14
		3.5.1. Correlated noise (?)	14
2. Instrument and data model	2	3.6. Data Selection and splits	14
2.1. Instrument overview	2	3.7. Mapmaking	14
2.2. Observation Strategy and Field Selection	3	4. Error Propagation, Simulations, and Null-tests	15
2.3. Data model	4	4.1. White Noise Simulations	15
2.4. Data overview	6	4.2. Transfer Functions	16
		4.3. Null-tests	17
3. COMAP Analysis Pipeline	8	5. Conclusions	17
3.1. Pipeline Overview	8		
3.2. Data Segmentation	9		
3.3. Filtering	10		
3.3.1. Normalization	10		
3.3.2. Removal of Pointing Templates	10		
3.3.3. Polynomial filter	10		
3.3.4. Principal Component Analysis (PCA) filter	10		
3.3.5. Masking	11		
3.3.6. Decimation	12		
3.4. Calibration	12		

1. INTRODUCTION

Cosmology ranks among the great success stories of modern science, and during the last few decades a solid theoretical framework for understanding early-universe physics has been established. As a concrete example, half a century ago physicists argued whether the Hubble constant was 50 or $100\ \text{km s}^{-1}\ \text{Mpc}^{-1}$. Today the discussion is whether it is 68 or $73\ \text{km s}^{-1}\ \text{Mpc}^{-1}$ —and whether that tiny difference may point toward ground-breaking new physics. The key to this success has been to map large-scale structures in the early universe, from small perturbations in the cosmic microwave

background (CMB) to the three-dimensional positions of current-day galaxies. This development is the direct result of the concerted efforts of thousands of scientists working for decades to gradually improve detector performance, data analysis methods, and theoretical models, and the feedback between these.

While containing a treasure trove of information regarding fundamental physical processes, cosmological density perturbations are also notoriously difficult to measure robustly due to their weak observational signatures. For instance, large-scale CMB temperature variations have a typical peak-to-peak variation of a few tens of microkelvins, while gravitational lensing induces a characteristic shear in the apparent ellipticity of individual galaxies of no more than a few percent. Detecting such weak signals in the presence of noise and systematic errors is a massive technological challenge, and requires both exquisite instrumentation and data analysis techniques.

As summarized by Kovetz et al. (2017), intensity mapping is a fairly recent development in this field, that holds the promise of surveying large cosmological volumes rapidly and with a relatively low cost. The fundamental idea behind this technique is to use redshifted line emission from molecular gas, for instance in the form of H_2 , CO, or C II, as a tracer for the underlying density field. Large volumes along a given line-of-sight may be surveyed simultaneously with a single spectrometer, and by scanning this spectrometer across the sky, a full 3D density map may be derived.

COMAP is a pathfinder intensity mapping experiment that targets carbon monoxide (CO) in the frequency range between 26 and 34 GHz, using a 19-element receiver mounted a 10.4 m single-dish telescope, and each detector is coupled to a 1024-channel ROACH spectrometer. Since the rotational $J = 1 \rightarrow 0$ transition of the CO molecule corresponds to a restframe emission frequency of 115.27 GHz, this means that COMAP should be sensitive to cosmological CO emission at a redshift of $z = 2.44\text{--}3.4$, with an angular resolution of $4'$ FWHM and a line-of-sight resolution of $\Delta z = 0.008$.

At the heart of the COMAP detector chain are cryogenically cooled HEMT low-noise amplifiers (LNA), which have excellent noise and systematics performance, which results in an overall system temperature of 35–45 K across the relevant frequency range. At the same time, the theoretically predicted signal from high-redshift CO emission is expected to be no more than a few microkelvins per COMAP resolution element (or “voxel”). Thus, the raw instrumental noise must be reduced by many orders of magnitude before a statistically significant detection may be achieved. In practice, this

Table 1: Frequency range of each COMAP sideband (SB).

Band	SB	Freq. (GHz)
A	LSB	26–28
	USB	28–30
B	LSB	30–32
	USB	32–34

is done by repeatedly observing the same part of the sky using multiple detectors, and thereby gradually increase the sensitivity per voxel. For this to succeed, however, it is necessary to suppress systematic contributions from atmospheric variations, sidelobe contamination, standing waves, Galactic foregrounds etc. by a corresponding amount.

The current paper describes the first-season COMAP data analysis pipeline, which aims to produce clean maps from raw time-ordered COMAP observations. This includes calibration, pointing reconstruction, data selection, filtering, and mapmaking. The rest of this paper is organized as follows: First, in order to establish useful notation and conventions, we give a brief introduction to the COMAP instrument in Sect. 2, while referring the interested reader to for full details. Next, we provide a high-level overview of the analysis pipeline in Sect. 3.1, before specifying each step in Sects. 3.3–3.7. Consistency tests are considered in Sects. 4–4.3, and we summarize and conclude in Sect. 5.

2. INSTRUMENT AND DATA MODEL

Before describing the COMAP analysis pipeline, we provide a brief overview of the instrument itself, and define an explicit data model. A more detailed description of the instrument can be found in a separate paper (in prep.).

2.1. Instrument overview

The COMAP instrument is located at Owens Valley Radio Observatory (OVRO) in California, USA, and saw first light on **January 1st, 2018**. It is mounted on a telescope that was originally built for the Millimeter Array at OVRO, then used as a part of the Combined Array for Research in Millimeter-wave Astronomy (CARMA) experiment, and has now been repurposed for COMAP. The primary and secondary mirrors of this telescope have diameters of 10.4 m and 1.1 m, respectively.

COMAP observes in the Ka band (26–34 GHz) with a $4'$ FWHM angular resolution. The COMAP receiver comprises 19 independent detector chains, each often informally referred to as a “feed”. Each feed consists of

the following elements, ordered according to their position in the signal path:

1. A feedhorn that collects incoming photons.
2. (Optionally) a polarizer that separates the radiation into circular two orthogonal polarization states; 15 feeds have a two-stage polarizer, 2 feeds have a one-stage polarizer, and 2 feeds have no polarizer.
3. An indium-phosphite low-noise amplifier (LNAs) integrated in the form of a monolithic microwave integrated circuit (MMIC) that incorporates six high-electron mobility transistors (HEMTs).
4. Two down converter modules (DCMs). DCM1 shifts the 26–34 GHz frequency band down to 2–10 GHz. DCM2 first splits the 8 GHz bandwidth into two bands A (2–6 GHz) and B (6–10 GHz), which are split again into two 2 GHz wide sidebands (SBs) each. The SBs are labelled by “lower” (LSB) or “upper” (USB) corresponding to the original frequency. Explicit frequency ranges for each SB is listed in Table 1.
5. Two CASPER “Roach2” FPGA-based spectrometers that splits each SB into 1024 separate 2-MHz wide frequency channels.

After detection, each spectrum is averaged in software into 256 8 MHz wide frequency channels, and recorded with a sample frequency of 50 Hz, resulting in a total data production rate of 13 GB/hour or 111 TB/year.

Different feeds have different polarizer solutions for development and testing purposes. As discussed later, one of the main systematic effects observed in the COMAP system are standing waves, i.e., electromagnetic waves reflected between the boundaries of various optical elements. The magnitude and impact of this effect depends sensitively on the polarization properties of the signal, and different solutions were therefore explored in the early phases of the experiment. An important finding from these tests was that feeds without a polarizer are significantly more noisy than the rest, and the two pixels in question are therefore excluded from the analysis. Similarly, a 20th pixel was added as a blind feed, to help distinguish internally and externally generated signals, and this is also excluded from the science analysis.

To achieve adequate sensitivity, the COMAP receiver is cryogenically cooled to 15–18 K within a cryostat that is situated between the primary and secondary mirror. To protect the instrument from environmental changes, a weather shield is installed on top of the cryostat. However, it was found during the operations that this

weather shield excited standing waves, and an improved design was installed in June 2020 that significantly reduces this effect.

To support both frequent and accurate gain estimation, COMAP employs a dedicated calibration vane that is directly attached to the cryostat housing. This vane is automatically moved in front of the feedhorn array at the beginning and end of each observation (each lasting for about one hour; see Sect. 2.2), fully filling the field of view of each pixel. The temperature of the calibration vane is monitored with thermometers, from which the responsivity of the system may be derived.

2.2. Observation Strategy and Field Selection

COMAP is a blind survey, observing three main areas of the sky. All fields have been chosen to overlap with The Hobby-Eberly Telescope Dark Energy eXperiment (HETDEX) (Hill et al. 2008), which targets Ly α emitting galaxies at $1.9 < z < 3.5$. These fields are selected to have faint foregrounds, and are possible to observe for several hours throughout the day. The three main fields are named CO2 (centered on $1^{\text{h}} 41^{\text{m}}$ RA 0° dec), CO6 (centered on $15^{\text{h}} 4^{\text{m}}$ RA, $+55^{\circ}$ dec), and CO7 (centered on $11^{\text{h}} 20^{\text{m}}$ RA, 52.5° dec). CO2 lies within the Spitzer/HETDEX Exploratory Large Area (SHELA; Papovich et al. 2016) which is in the fall field of HETDEX, along the SDSS Stripe 82, while CO6 and CO7 both lie within the HETDEX spring field.

In addition to the main science fields, we observe several Galactic fields for ancillary science applications, in particular to constrain anomalous microwave emission (AME) or spinning dust emission around 30 GHz. These fields include the Andromeda galaxy (FG4), the star forming region W43 (FG6) and a region within the Perseus constellation (FG7). We are also conducting a larger Galactic plane survey centered on FG6. The analysis of the galactic fields can be found in a separate paper (Harper et al. 2021).

To facilitate calibration with astrophysical sources, we also observe a handful of point sources, including Jupiter, the supernova remnants Taurus A (TauA) and Cassiopeia A (CasA), and the radio galaxy Cygnus A (CygA). Analysis of these observations is discussed in Sect. 3.4. Considering the COMAP beam size, Jupiter is the only true point source, and the others are also included in the Galactic analysis.

Table 2 lists all fields (excluding Jupiter) and their celestial coordinates. In Figure 1 we plot the elevation of each fields as a function of time during January 1st 2020, indicating when the fields are available for observation. Throughout the year, the individual position of

Patch	RA	dec
CO2	01:41:44.4	+00:00:00.0
CO6	15:04:00.0	+55:00:00.0
CO7	11:20:00.0	+52:30:00.0
FG4	00:42:44.4	+41:16:08.6
FG6	18:44:56.4	−01:59:52.0
FG7	03:40:49.0	+31:54:38.0
TauA	05:34:31.9	+22:00:52.2
CasA	23:23:24.0	+58:48:54.0
CygA	19:59:28.4	+40:44:02.1

Table 2: Celestial coordinates of all COMAP fields.

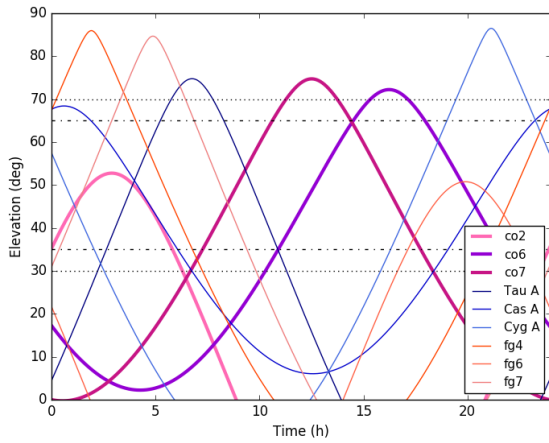


Figure 1: Elevation of CO (pink/purple), galactic (orange) and calibration (blue) fields as observed from Owen’s Valley Radio Observatory during January 1st 2020.

these peaks shifts in time, but the overall picture remains fixed.

The instrument is scanning the sky in two main patterns, constant elevation scans (CES) and Lissajous scans. The two observation strategies are alternated on a daily basis. CESs correspond to drift scans that are constant in elevation, whereas Lissajous scans follow an harmonic motion described by

$$Az = A \sin(at + \phi); \quad El = B \sin(bt), \quad (1)$$

where A, B are amplitude parameters that defines the size of the field, the ratio a/b determines the shape of the curve, and ϕ is a phase parameter. For a circular scan, we have $A = B$, $a = b$, and $\phi = \pi/2$. For COMAP, the ratio a/b is randomized between each scan. When the field has drifted through the telescope pointing, which typically takes 5–10 minutes, the telescope is repointed to the leading edge of the field. An example of the scan-

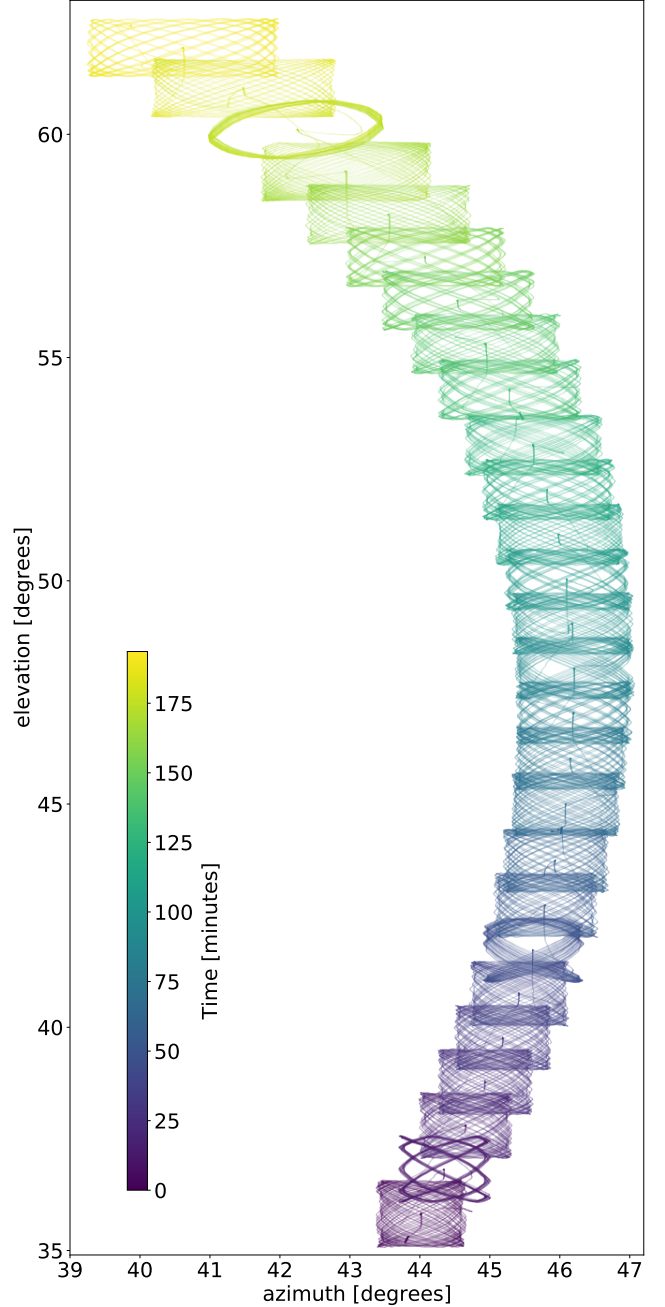


Figure 2: Movement of the telescope in azimuth and elevation for three consecutive observations (27 scans, 194 minutes) of the same field, with the Lissajous scanning strategy.

ning path for 3 hours of continuous observations is shown in Figure 2.

2.3. Data model

As described by ?, the COMAP detector readout for a single frequency channel may be modelled as

$$P_{\text{out}} = k_B G \Delta \nu T_{\text{sys}}, \quad (2)$$

where G is the gain, $\Delta\nu$ is the bandwidth, and T_{sys} is the system temperature of the instrument. The system temperature may be further approximated as

$$T_{\text{sys}} = T_{\text{receiver}} + T_{\text{atmosphere}} + T_{\text{ground}} + T_{\text{CMB}} + T_{\text{foregrounds}} + T_{\text{CO}}, \quad (3)$$

where T_{receiver} is the effective noise temperature of the receiver, $T_{\text{atmosphere}}$ is the noise contribution from the atmosphere, T_{ground} is ground pickup from far sidelobes, T_{CMB} is the contribution from the CMB, $T_{\text{foregrounds}}$ are continuum foregrounds (typically from the galaxy), and T_{CO} is the line emission signal from extragalactic CO, which is the main scientific target of the COMAP instrument.

To understand the challenge involved in measuring the cosmological CO signal, it is instructive to consider the order of magnitude and stability of each term in Eq. 3. By far the largest term is the receiver temperature, which is about 30 K for COMAP. However, since the instrument is cryogenically cooled, this term is relatively stable in time, and changes may be monitored through thermometers.

The second largest term is the atmosphere contributions, which typically contributes with 6–8 K. This term varies significantly on all time scales longer than a few seconds, and depends on external conditions including humidity, cloud coverage, ambient temperature etc. It is also strongly correlated between detectors and frequencies, since all feeds observe through essentially the same atmospheric column density at any given time; fortunately, the phase structure of the atmospheric fluctuations are uncorrelated on long time scales.

Next, ground pickup typically accounts for 2–4 K, and this term can be particularly problematic because it depends sensitively on the instrument pointing: If a sidelobe happens to straddle a strong signal gradient, such as the horizon or the Sun, several mK variations may be measured on very short time-scales and with an time-dependency that appears nearly sky synchronous.

The fourth term represents the CMB temperature of 2.7 K, which is both isotropic and stationary, while the fifth term represents astrophysical foregrounds of at most 1 mK, for instance synchrotron, free-free, and dust emission. Although these are sky synchronous, and in principle could confuse potential CO measurements, they also have very smooth frequency spectra, and are therefore relatively easy to distinguish from the cosmological CO signal, which varies rapidly with frequency. An important potential exception is line emission from other molecules than CO that happen to be located at other redshifts than those targeted by COMAP. However, as discussed by ?, there are no known molecules

with a sufficient cosmological abundance to represent a significant contaminant in the frequency range between 26 and 34 GHz. Finally, the cosmological CO line emission signal is expected to account for $\mathcal{O}(1\mu\text{K})$.

Whether it is possible to detect such a weak signal depends directly on the stability and sensitivity of the instrument. In this respect, the fundamental quantity of interest is the overall noise level of the experiment, which is dominated by random thermal Johnson noise caused by thermal motions of electrons within the electronics. The magnitude of these random thermal fluctuations is proportional to T_{sys} , with a standard deviation that given by the so-called radiometer equation,

$$\sigma_N = \frac{T_{\text{sys}}}{\sqrt{\Delta\nu \tau}}, \quad (4)$$

where τ is the integration time. Thus, since both the system temperature and the bandwidth are essentially fixed experimental parameters, the only way of reducing the total uncertainty is by increasing the integration time. As a concrete and relevant example, we note that an integration time of 70 hours is required to achieve a standard deviation of $1\mu\text{K}$ with a system temperature of 45 K and a bandwidth of 8 GHz.

In addition to the thermal and uncorrelated noise described by the radiometer equation, there are three main sources of correlated noise, namely gain fluctuations in the low-noise amplifiers, atmospheric temperature fluctuations, and time-dependent standing waves. All of these are expected to have a roughly $1/f$ -type spectrum, although with different particular properties. The fact that these sources of correlated noise are also strongly correlated between frequencies is very useful in order to filter out this noise in the analysis.

Equation 17 describes the detector output at any given time. To connect this to the actual measurements recorded by the detector, we adopt the following useful data model,

$$d_\nu^i(t) = \langle d_\nu^i \rangle (1 + \delta_G^i) [P_{\text{cel}}^i (s_{\text{cont}} + s_{\text{CO}}^\nu) + P_{\text{tel}}^i s_{\text{ground}} + n_{\text{corr}} + n_{\text{w}}^{\nu i}]. \quad (5)$$

Here $d_\nu^i(t)$ denotes the raw data recorded at time t for frequency channel ν in feed i ; $\langle d_\nu^i \rangle$ represents the corresponding time average; δ_G^i denotes feed dependent gain fluctuations; P_{cel}^i and P_{tel}^i are pointing matrices in celestial and telescope coordinate systems, respectively; s_{cont} denotes the celestial continuum sources, mainly from the CMB and galactic foregrounds; s_{CO}^ν is the CO line emission; s_{ground} is the ground signal picked up by the far sidelobes; and n_{corr} is the correlated noise component, mostly consisting of atmosphere fluctuations and standing waves. Factors with no feed or frequency index are

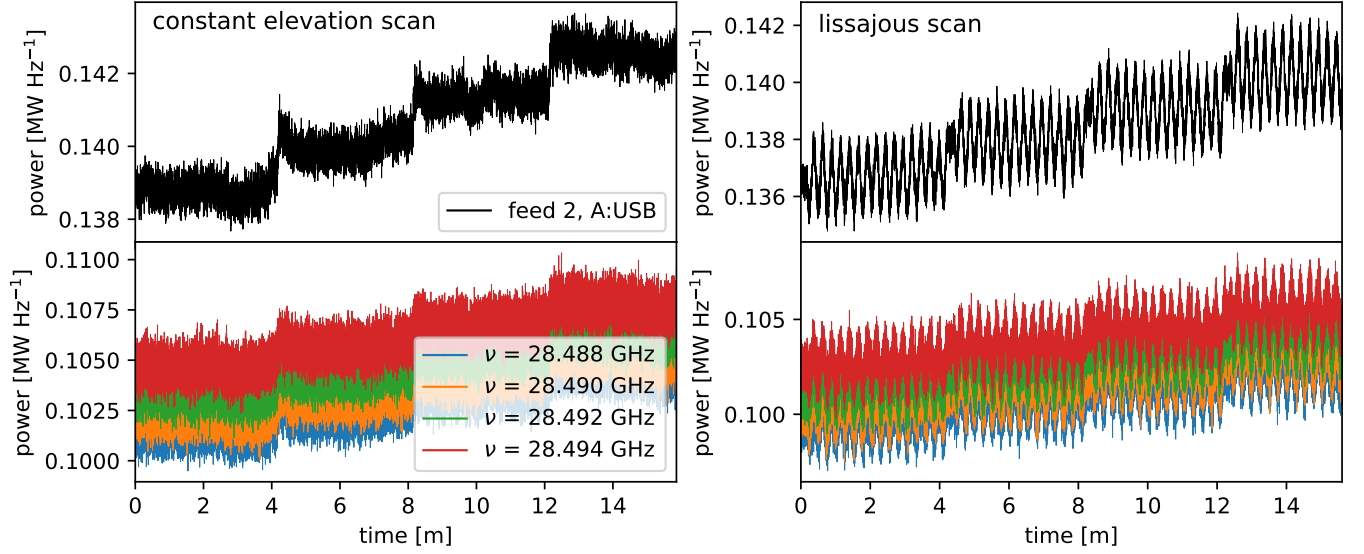


Figure 3: Raw data from the COMAP instrument. Here we see sideband averaged data (top) and examples of data from individual frequency channels (bottom). We show data from an observation using CES (left) and an observation using lissajous scans (right).

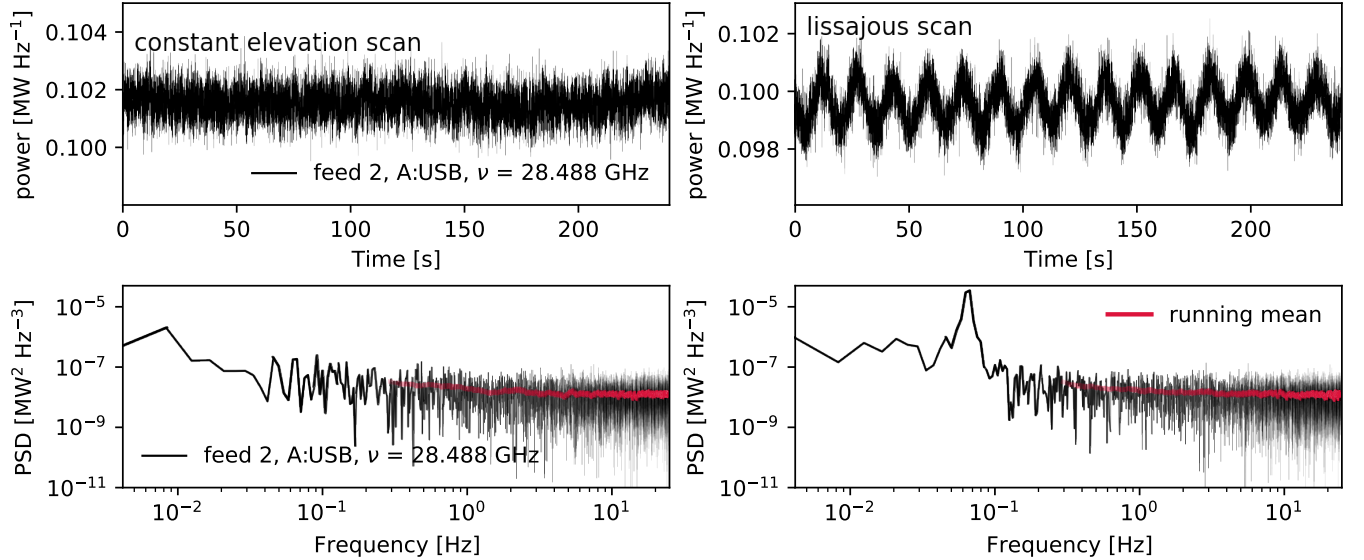


Figure 4: Raw data from an individual frequency channel of the COMAP instrument. Power is shown as a function of time (top), and the corresponding power spectral density (PSD) is also shown (bottom). We show data from a CES scan (left) and a lissajous scan (right).

assumed to be independent between frequencies (or at least strongly correlated), while factors with a ν label indicate parts of the model that are assumed to have non-smooth frequency dependence. The main purpose of the COMAP analysis pipeline is to characterize s_{CO}^ν given $d_\nu^i(t)$.

2.4. Data overview

Before presenting the analysis pipeline, we provide a preview of the raw time-ordered data (TOD) generated by the COMAP instrument, with the goal of building intuition that will be useful for understanding the purpose of each component of the analysis pipeline described in this paper. Figures 3 and 4 show examples of such raw TOD from the instrument using the CES (left column) and Lissajous (right column) scanning strategies. Perhaps the most obvious features in these plots are step-

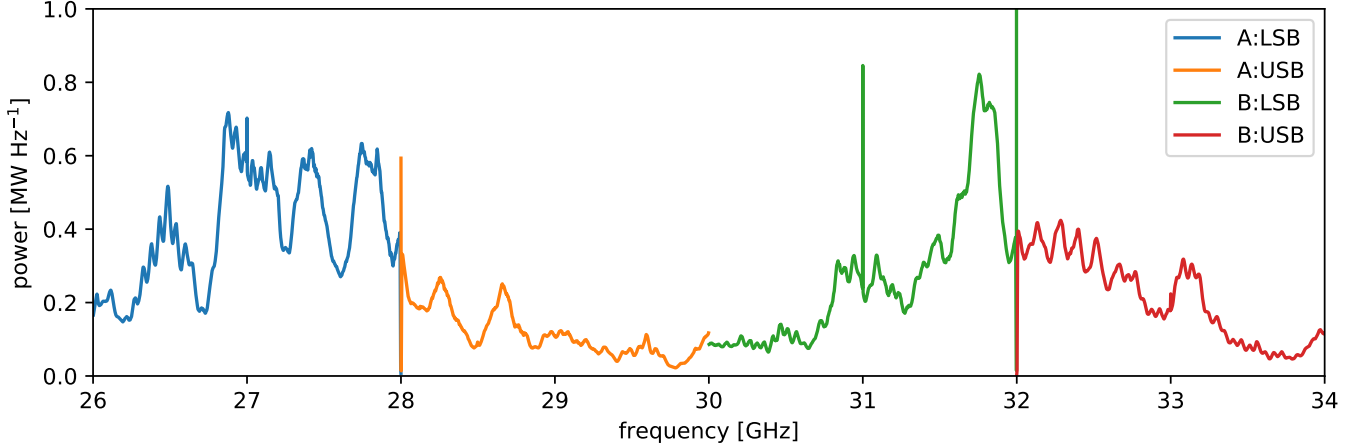


Figure 5: Time averaged raw data from each frequency channel on single feed of the COMAP instrument.

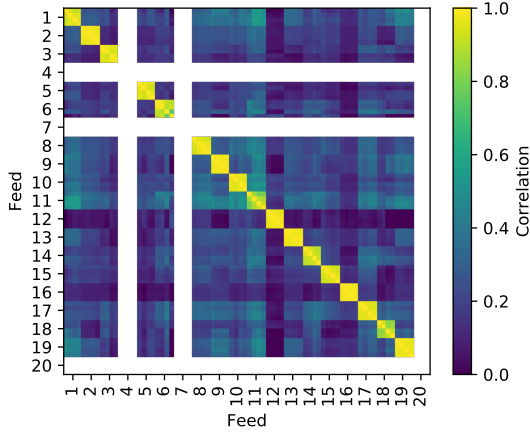


Figure 6: Correlation between the sideband-averaged data from the 20 feeds of the COMAP instrument, after pointing template subtraction.

wise changes in power as the telescope changes elevation during repointings between scans; see Sect. 2.2. The Lissajous scans additionally show oscillations in power as the telescope changes elevation during the scan, since the telescope looks through a thicker slab of atmosphere at lower elevations, and this increases the atmospheric contribution to the system temperature.

The top panels in Fig. 4 show an individual frequency channel for a single scan (i.e., stationary observation period), while the bottom panel shows the corresponding power spectral density (PSD). For the CES case, the PSD is relatively featureless, with an overall shape that looks consistent with a typical $1/f$ noise spectrum. For the Lissajous case, an additional strong peak is seen around 0.007 Hz, which matches the scanning period of 14 sec, and this corresponds to the periodic atmospheric variations seen in the panels above.

Figure 5 shows the time averaged data for all frequency channels of a single feed. The spectral shape is mostly determined by the average gain as a function of frequency, and as such is a characteristic of the low noise amplifier attached to each feed. This average gain is a pure instrumental effect, and not associated with the true sky signal, and therefore simply corresponds to a normalization factor that should be calibrated out before higher-level analysis. However, some of the spectral shape is also determined by the fact that the system temperature also changes with frequency, and in some cases results in large spikes within specific frequency ranges. Separating these two effects is a main goal of the calibration procedures described below.

In Fig. 6 we plot the correlation,

$$C_{ij} = \frac{\langle \hat{d}^i \hat{d}^j \rangle}{\sqrt{\langle \hat{d}^i \hat{d}^i \rangle \langle \hat{d}^j \hat{d}^j \rangle}}, \quad (6)$$

between the power, \hat{d}^i recorded by any two radiometers, i and j , after subtracting both the mean and an elevation-dependent template of the form

$$T_{\text{el}}(t) \propto \frac{1}{\sin \text{el}(t)} \quad (7)$$

from each TOD segment, and averaging over all frequencies within each sideband for each radiometer. Here we first note that the data from different sidebands of the same feed are strongly correlated. This is because both main sources of correlated noise in the COMAP data, namely gain fluctuations and atmospheric fluctuations, are common for sidebands within a given pixel. In contrast, sidebands for different feeds share the atmospheric fluctuations (and also some standing waves), but have mostly independent gain fluctuations, and this results in

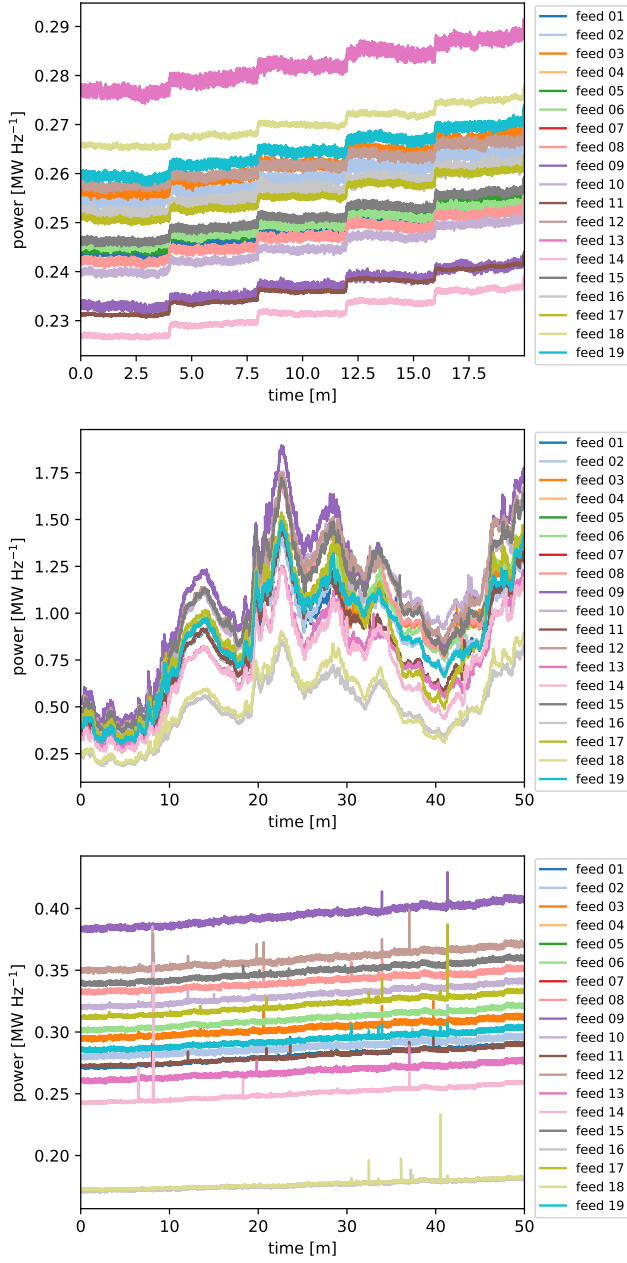


Figure 7: Feed averaged COMAP TOD recorded under various observing conditions. The top panel shows data observed under normal conditions, and is dominated by instrumental noise. The middle panel shows data observed under poor weather conditions with a thick cloud coverage, resulting in large coherent power fluctuations observed by all feeds. This third panel shows data with strong spikes, which may for instance happen during periods with high insect activity.

lower overall correlations, but still typically in the 10–40% range. Accounting for and mitigating such correla-

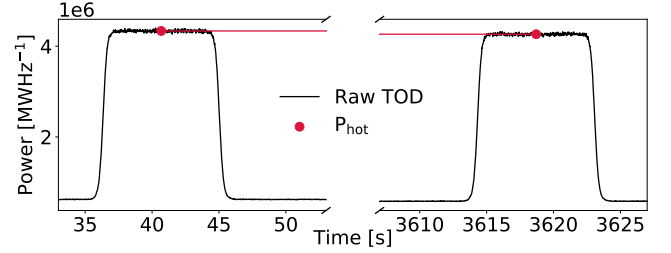


Figure 8: Zoom-in on calibration vane observations made at the beginning and end of each observation period, typically lasting for about one hour.

tions will clearly be essential in order to extract robust science from these observations.

The quality of the COMAP data depends strongly on the observation conditions at any given time, as illustrated in Fig. 7. The top panel shows an observation made under normal conditions, while the middle panel shows an observation made under poor weather conditions with thick cloud coverage. The bottom panel shows a data segment with strong spikes, which may for instance be caused by insects flying in front of the focal plane. Automatic identification and removal of such features is clearly an important and necessary component of the pipeline.

Finally, Fig. 8 shows a zoom-in of the calibration vane observations that are made at the beginning and end of each observation period. Since the ambient temperature is about one order of magnitude higher than T_{sys} , the measured power is also correspondingly about one order of magnitude higher, and this bright and known signal allows for a clean calibration estimate. Note that, unless specified otherwise, these data segments are always removed prior to data analysis, as they would otherwise compromise any filtering that may be applied to the data.

3. COMAP ANALYSIS PIPELINE

3.1. Pipeline Overview

We are now ready to present the COMAP analysis pipeline, which is designed to process the raw data discussed in Sect. 3.1 into calibrated and cleaned CO maps. The main steps of this pipeline are schematically illustrated in Fig. 9.

The processing starts with raw Level-1 files, which contain raw data as recorded by the instrument, together with pointing information and house-keeping data. These files typically each contain about one hour of observations, including calibration vane observations at the beginning and end, as well as a corresponding estimate of the system temperature T_{sys} based on the same calibration measurements. We denote each one hour

of data as one observation, and assign it an individual observation IDs (abbreviated ObsID). Each observation consists of several scans, where one scan is the observation done between two re-pointings of the telescope. During this time the telescope performs the same motions around a fixed point in azimuth and elevation while the source drifts through, and the instrumental properties are consequently assumed to be stationary within each scan. The entry denoted `scan_detect` in Fig. 9 indicates a dedicated code that partitions each observation into individual scans, based on pointing information, and organizes these in an observation database.

The main processing takes part in `l2gen`, which generates calibrated and cleaned TOD, and stores these in so-called Level-2 files. This is achieved through a series of filters (see Sect. 3.3), and applying a time-varying gain normalization (see Sect. 3.4). This stage also evaluates basic goodness-of-fit statistics, and defines a basic frequency mask that excludes missing or broken data for the current scan. Finally, this step also performs data averaging, reducing the spectral resolution of the data according to user input; in most analyses, we typically reduce the resolution from 2 MHz to 32 MHz, resulting in a computational speed-up of subsequent steps of a factor of 16, and a memory saving for storing final maps of a factor of 256.

Next, the `accept_mod` program reads in the goodness-of-fit statistics and basic frequency mask produced by `l2gen`, and produces a list of accepted observations as defined by user-specified thresholds for each statistic. Examples of relevant statistics used for this purpose are χ^2 per observation, f_{knee} , or Solar distance. The output from this process is called an *accept list*, which also is used to define data splits.

Converting time-ordered data into pixel-ordered data is done by a mapmaker called `tod2comap` (see Sect. 3.7). As shown in the following, the adopted filters result in very nearly uncorrelated white noise, and the current implementation of `tod2comap` accordingly adopts simple binning into voxels. Finally, from these maps we can extract an estimate of the CO power spectrum, or the voxel intensity distribution, or perform general astrophysics and cosmology inference, all of which happens in `comap2ps`.

3.2. Data Segmentation

As described above, we define a *scan* to be the time of observation between re-pointings of the telescope. During a scan the telescope performs the same movements (e.g., constant elevation or Lissajous) centered on the same azimuth and elevation positions for 3 to 10 min-

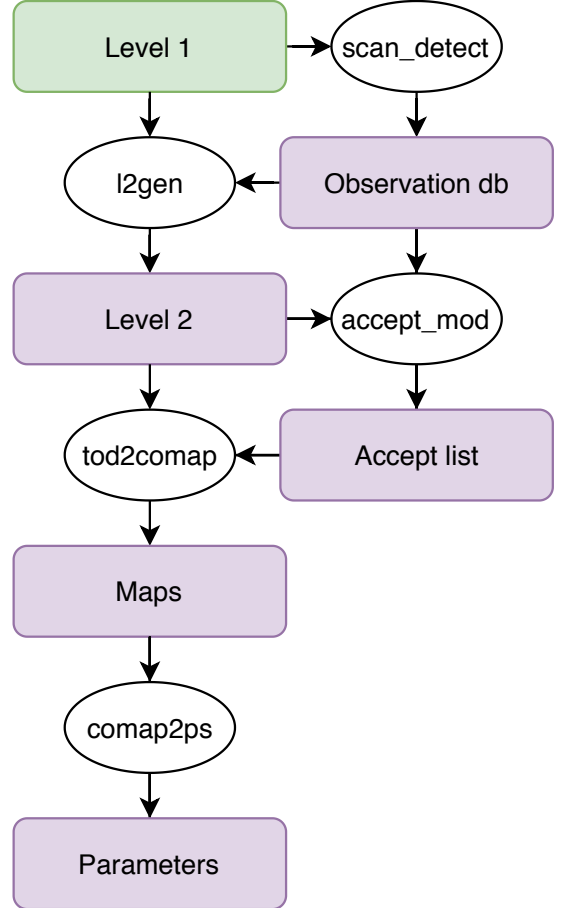


Figure 9: Flow diagram of the analysis pipeline. The purple boxes are products of this pipeline, whereas the green is the input data (raw data from the telescope and some housekeeping data). The white ellipses represents different modules of the code.

utes while the observed source or field drifts through. In the subsequent analysis each scan is analyzed independently, and is converted into an individual Level-2 file.

The purpose of the `scan_detect` code is to identify all scans within all observation periods, and sort these according to astrophysical target. This selection is based on telescope house-keeping information, and in particular on an observation flag that records the state of the telescope at any given time. However, recognizing that human errors do happen, automatic simple sanity tests are also performed, for instance to validate that the specified field name agrees with the recorded pointing.

The basic result from this selection process is called a *runlist*, which simply is a list of obsIDs sorted according to source. For each obsid, we list all scans within that obsid, including basic information such as the modified Julian date (MJD) of the start and end of the scan,

as well as the scanning mode (e.g. circular or constant elevation) and mean pointing information. The calibration vane observations are defined to be the first and last scans within each obsID, but flagged as calibration measurements.

3.3. Filtering

As described in Sect. 3.1, the COMAP TOD exhibit a wide range of non-CO-related contributions, both of instrumental and external origin. These must be suppressed by orders of magnitude prior to map-making in order to extract the astrophysically valuable signal. With this goal in mind, we introduce four specific filters, each targeting one class of artefacts.

3.3.1. Normalization

The first filtering operation we introduce is data normalization. This is done simply by dividing the raw TOD, $P_{\text{out},i}$, by its own running mean, and then subtracting the overall mean,

$$D(\nu, t) = \frac{P_{\text{out}}(\nu, t)}{\langle P_{\text{out}}(\nu, t) \rangle} - 1. \quad (8)$$

Here t is a time sample index, whereas ν denotes frequency channel. This operation is performed separately on each frequency channel.

The main purpose of this step is to equalize (i.e., "flatten") the instrumental bandpass, as illustrated in Fig. 5, and effectively establish data with appropriate relative calibration. The main practical advantage of doing so is that the amplitude of common-mode contaminants, such as gain-induced correlated noise or atmospheric fluctuations, become comparable across all frequencies within a single sideband, and therefore much easier to filter out. The same also holds true for broad-band astrophysical contributions, such as the CMB or foregrounds, which also must be removed prior to CMB extraction.

Note also that with the definition in Eq. 8, the noise level of $D(\nu, t)$ is given by the sample rate and bandwidth alone in the ideal case, and should equal $1/\sqrt{\tau\Delta\nu}$. Calibration into physical units is performed simply by multiplying $D(\nu, t)$ with T_{sys} . $D(\nu, t)$ is as such a particularly convenient function for goodness-of-fit tests, and it will serve as our main object of interest in the following.

3.3.2. Removal of Pointing Templates

As the telescope changes elevation we are looking through different amounts of atmosphere, leading to changes in the power received by the detectors. This effect can be modeled by a simple expression for the optical depth of the atmosphere, τ

$$\tau(\text{el}) = \frac{\tau_0}{\sin(\text{el})}, \quad (9)$$

where τ_0 is the optical depth of the atmosphere at zenith, and el is the elevation.

We also sometimes see structure that is correlated with the azimuthal position of the telescope, presumably mostly because of ground pickup by far sidelobes.

We want to remove the fluctuations in our data that only result from these changes in elevation and azimuth. We do this for each frequency channel separately by fitting two following pointing templates to the timestream. Specifically we assume that

$$d = \frac{g}{\sin(\text{el}(t))} + a \text{az}(t) + c + n, \quad (10)$$

where g , a and c are constants and n is some Gaussian noise with constant variance. We then find the best fit values (i.e. least squares) of g , a and c , and use the values of g and a to remove the pointing templates from the timestream:

$$d_{\text{after}} = d_{\text{before}} - \frac{g}{\sin(\text{el})} - a \text{az} - \left\langle \frac{g}{\sin(\text{el})} + a \text{az} \right\rangle, \quad (11)$$

where $\langle \rangle$ denotes the mean value in time for a specific frequency channel. See figure ??.

For longer scans we divide the data into different segments of roughly 4 minutes each, and perform this template fit and removal separately on each of these segments of data.

3.3.3. Polynomial filter

For each timestep, we fit a low-order polynomial to the data in frequency space for each of the sidebands, and subtract it. This removes all common modes over all frequencies. The goal of this is to remove the $1/f$ noise specific to each detector, and to remove any continuous foregrounds. As the foreground are expected to be smooth over all frequency bands, they should be effectively removed by the polyfilter. It will, however, also remove some of the largest modes of the CO signal. These modes correspond to a physical scale of around 100 Mpc/h, and are generally not targeted by COMAP.

3.3.4. Principal Component Analysis (PCA) filter

Some systematics, like standing waves or atmospheric variations have the same time-variations for many or most feeds. In these cases we can do a PCA analysis on the whole dataset to fit and remove these common modes. Intuitively this means finding the functions of time which explain the largest amount of the variance between the different frequencies across all the different feeds. These functions are what we call the leading PCA components.

Consider all the data corresponding to one scan as a data matrix D , where each row is the TOD corresponding to a single frequency channel on a single feed.

Thus D is a matrix with dimensions $n_{\text{freq}} \times n_{\text{samp}}$, where $n_{\text{freq}} = n_{\text{feeds}} \cdot n_{\text{sidebands}} \cdot n_{\text{freq per sideband}} = 19 \cdot 4 \cdot 1024$ is the total number of frequencies added up from all sidebands and feeds, and n_{samp} is the number of samples in time

$$D = \begin{bmatrix} D_{11} & \dots & D_{1n_{\text{samp}}} \\ \vdots & \ddots & \vdots \\ D_{n_{\text{freq}}1} & \dots & D_{n_{\text{freq}}n_{\text{samp}}} \end{bmatrix}. \quad (12)$$

Consider the covariance matrix, C , of this data

$$C = D^T D. \quad (13)$$

The eigenvectors, \mathbf{v}_k , of this matrix corresponding to the highest eigenvalues are exactly the PCA components we are looking for. We use an iterative method to extract the first few leading components.

For each frequency (on each feed) we find the amplitudes of the PCA components by finding the overlap between the data, \mathbf{d} , and the component \mathbf{v}_k

$$a_k = \mathbf{d} \cdot \mathbf{v}_k = \sum_{i=1}^{n_{\text{samp}}} d_i v_k^i, \quad (14)$$

where \mathbf{d} is the timestream for the given frequency channel. The leading PCA component are then subtracted from the data

$$\mathbf{d}_{\text{after}} = \mathbf{d}_{\text{before}} - \sum_{i=1}^{n_{\text{comp}}} a_i \mathbf{v}_i, \quad (15)$$

where n_{comp} is the number of leading components removed (typically four).

Figure 10 shows the three leading PCA components for a typical scan.

3.3.5. Masking

Some times individual frequency channels or groups of nearby frequency channels show some deviation from the expected uncorrelated white noise behaviour, even after applying all the filters. This could manifest in an excess noise that is correlated in time, or in correlations between different frequency channels. We want to mask out these frequency channels so that the contribution from these “bad” channels are not mixed with the otherwise good ones.

To determine which frequencies should be masked we first perform the poly- and PCA filters on a copy of the dataset.

We then use two main approaches to identifying individual or groups of frequency channels to be masked.

The first approach uses the fact that the expected correlation between two independent Gaussian variables

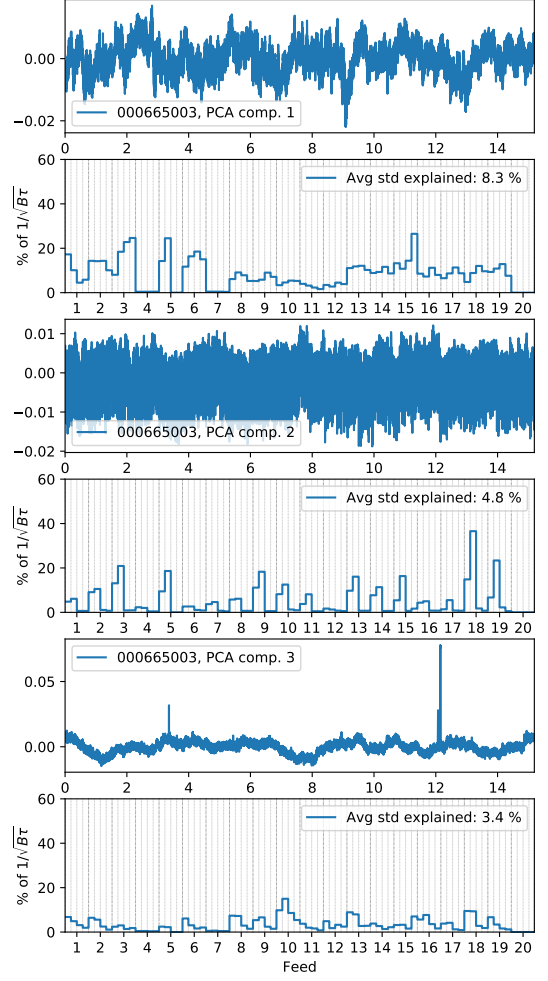


Figure 10: The three leading PCA components of a typical scan, and the amount of standard deviation it explains.

(for large n_{samp}) is given by $1/\sqrt{n_{\text{samp}}}$, where n_{samp} is the number of samples the correlation is calculated on.

This means that, after subtracting the expected correlation induced by the polyfilter, we know the statistics describing *good* data, and can identify bad data as deviations from these statistics. We look at entries within squares of different sizes as well as sets of columns within the correlation matrix and see if the average of the absolute value of the correlations within this region deviates from the values expected, and mask out the corresponding channels if the deviations are outside the acceptable limits.

The second approach is to calculate a set of diagnostics for individual frequency channels (like the average correlation of the channel to all the others in that band or the average of the absolute value of the same). We can then compare the values of these diagnostics for the different channels and look for large outliers.

In addition to these approaches we look specifically for *edge correlations*, that is, correlations between individual frequencies at the edge of each sideband with the corresponding frequencies at the edge of another sideband. This is because we know that some aliasing between these frequencies can occur, and we want to remove frequencies where this aliasing is large.

After this mask has been found, we go back to the original dataset and do all the filtering and analysis again, but now only on the data that is not masked.

3.3.6. Decimation

All the analysis up until now have been performed on the full frequency grid (1024 channels per sideband). For mapmaking purposes, however, we typically don't need this kind of resolution, so we want to co-add several frequency channels together to a single low resolution channel.

$$d_i^{\text{lowres}} = \frac{1}{\sum_m w_m} \sum_{m=(i-1)n_{\text{dec}}+1}^{in_{\text{dec}}} w_m d_m^{\text{highres}}, \quad (16)$$

where d_i^{lowres} is the timestream of frequency channel i in the low resolution frequency grid, d_m^{highres} is the timestream of frequency channel m in the high resolution frequency grid, n_{dec} is the number (usually 16) of high resolution frequencies to be combined in each low resolution frequency channel and $w_m = 1/\sigma_m^2$ is the inverse variance of the timestream in frequency channel m of the high resolution data (w_m is zero for masked frequencies).

3.4. Calibration

Revisiting the equations from 2.3. The detector readout can be written as

$$P_{\text{out}} = k_B G \Delta \nu T_{\text{sys}}, \quad (17)$$

where G is the gain, $\Delta \nu$ is the bandwidth, and T_{sys} is the system temperature. T_{sys} is a fundamental parameter for characterizing the performance of a radio telescope, and should be as low as possible. The noise level of the data is proportional to T_{sys} , which is the added noise from the system, with contributions from both the sky and the receiver. The standard deviation of the noise is given by the radiometer equation

$$\sigma_N = \frac{T_{\text{sys}}}{\sqrt{\Delta \nu \tau}}, \quad (18)$$

where τ is the integration time. One way of estimating the T_{sys} is to place a hot load of known temperature in front of the telescope.

3.4.1. Calibration using a reference load

Ideally we would put a load of a known temperature in front of the telescope and above the atmosphere and compare the measured output power with the output power measured with no load. A good approximation is to use an ambient temperature load that covers the receiver feedhorn. Assuming that the telescope, ground spillover, and the atmosphere have the same temperature as the ambient load, the output power will be the same as if the load was above the atmosphere. Taking into account the vertical temperature profile and the distribution of the absorbing components in the atmosphere, the corrections are only a few percent for the relevant wavelengths. To measure the system temperature we compare the readout when we have a vane P_{amb} and when we look at the cold sky P_{cold} . From eq. (17) we can estimate T_{sys} as

$$T_{\text{sys}} = \frac{T_{\text{amb}} - T_{\text{CMB}}}{P_{\text{amb}}/P_{\text{cold}} - 1}, \quad (19)$$

where T_{amb} is the ambient temperature and T_{CMB} is the cold sky temperature. We then multiply the data $D(\nu, t)$ with the T_{sys} measurement to go from (normalized) detector units to K

$$D(\nu, t) \Big|_{\text{K}} = D(\nu, t) \Big|_{\text{adu}} \langle T_{\text{sys}}(\nu, t) \rangle. \quad (20)$$

As both the atmosphere and the receiver gain varies over time, the measurements of T_{sys} vary over time. To get the most accurate estimation, we make use of the vane at the beginning and end of each observation.

3.4.2. Photometric Calibration

Another way of calibrating the signal is to use an astrophysical source with a known temperature. This way we can also check if our estimations of T_{sys} with the vane are correct. A good source is Jupiter, but Tau A has also been used. As the angular size of Jupiter Ω_J differs from the solid angle of our beam Ω_{beam} , we have to take that into account. The expected temperature is given by

$$T_{\text{exp}} = T_J \frac{\Omega_J}{\Omega_{\text{beam}}}, \quad (21)$$

where T_J is the brightness temperature of Jupiter. We use a model for the T_J taken from measurements made by the WMAP experiment (Weiland et al. 2011). Performing a simple weighted power-law fit to the measurements gives the following relationship

$$\log(T_J) = 0.149 \log\left(\frac{\nu}{22.8 \text{ GHz}}\right) + 2.15 \text{ K}. \quad (22)$$

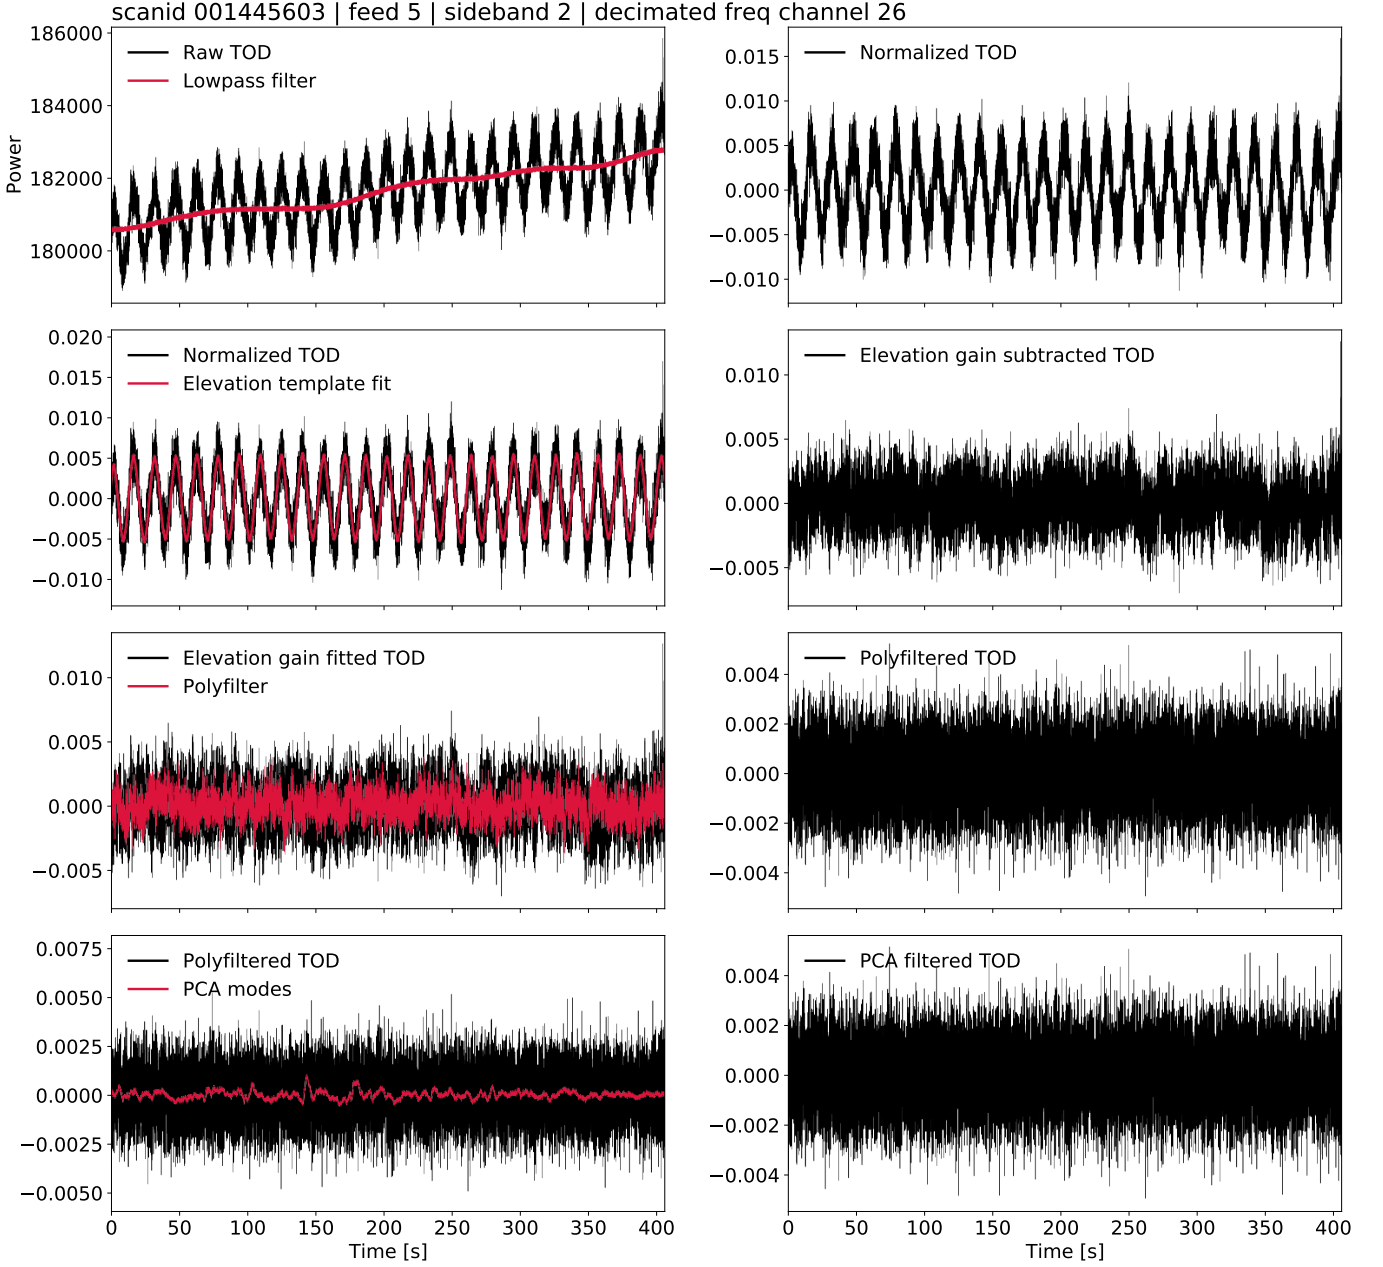


Figure 11: How the data changes during the course of the pipeline. Each row shows the data before (left figure) and after (right figure) subsequent steps in the pipeline. From top to bottom, the filters shown are 1. Normalization. 2. Elevation gain subtraction. 3. Polyfilter. 4. PCA filter.

The solid angle of Jupiter depends on the current distance r between Earth and Jupiter. Based on the reference value $\Omega_{5.2\text{AU}} = 2.481 \times 10^{-8}$ sr, we can estimate Ω_J to be

$$\Omega_J(r) = 2.841 \times 10^{-8} \left(\frac{5.2 \text{ AU}}{r} \right)^2. \quad (23)$$

The COMAP beam model is

$$\Omega_{\text{beam}} = \left[0.2593 \ln \left(\frac{\nu}{1 \text{ GHz}} \right)^2 - 1.965 \ln \left(\frac{\nu}{1 \text{ GHz}} \right) + 3.865 \right] \times 10^{-5} \text{ sr}. \quad (24)$$

However, as the receiver is not absorbing all of the incoming radiation, the temperature measured by the receiver, the antenna temperature T_{ant} , is lower than T_{exp}

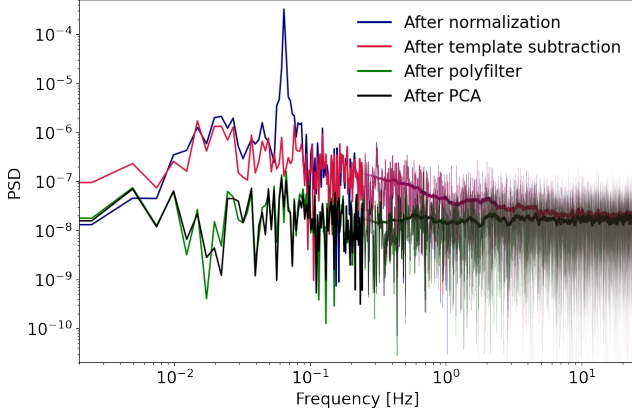


Figure 12: Effect of subsequent pipeline filters on the signal power spectral density. Towards higher frequencies, the thickness of the lines are gradually reduced, and the running mean of each signal is overplotted. The running mean of the blue and red, as well as the green and black lines are indistinguishable.

by a factor η

$$T_{\text{ant}} = \eta T_{\text{ext}}. \quad (25)$$

η is called the aperture efficiency and is another measure of the sensitivity of a radio telescope. η is the ratio of power from the source as measures by the receiver to the total power from the source hitting the telescope. The usual way of measuring η is to point the telescope at a point source with known surface temperature, like Jupiter.

3.5. Noise characterization

[Jonas] Necessary for the correct weighing for the mapmaking. The noise spectrum is assumed to be

$$N(\nu) = \sigma_0^2 \left(1 + \left(\frac{\nu}{f_{\text{knee}}} \right)^\alpha \right), \quad (26)$$

where σ_0^2 is the white noise level, whereas α and f_{knee} characterize the non-white noise. σ_0^2 is estimated by calculating the variance between two neighbouring samples for a given frequency and feed. α is the slope of the $1/f$ noise, and f_{knee} is the frequency where the lines of the $1/f$ noise and the white noise intersect. As for now, these values are set to $\alpha = 2.0$ and $f_{\text{knee}} = 10^{-6}$. Figure xx shows how the noise changes during the processing (?).

3.5.1. Correlated noise (?)

3.6. Data Selection and splits

We go through several stages of data selection in the COMAP pipeline. First of all, as discussed in section !!!!!, in each individual scan, we mask out individual

and groups of frequencies that behave badly, to keep them from contaminating the rest of the data.

The main data selection step, however, happens after we have already generated level 2 files. In *accept mod* we first go through all the data to make a database of various statistics and housekeeping data for each sideband of each feed and each scan. This database gives us a powerful way to characterize all the data and find patterns and correlations.

The next step is to use the statistics database to impose cuts on the data. For example we can exclude all the data with an elevation lower than 35 degrees, or any sideband where there spikes in the raw data for that particular scan. This procedure produces an accept list that stores whether or not a given sideband of a give scan is accepted or rejected.

The fact that this main data selection step happens after the main filtering steps, which are the most computationally expensive part of the pipeline, makes it much easier to change the cuts or just compare results using different cuts.

With the statistics database and the accept list in hand, we have all we need to split the data in any way we would want. We could for example do a half mission split, where the chronologically first half of the accepted data is split from the second half, or we could split the observed data according to the local sidereal time, or according to any function of the statistics we have gathered. These splits are very useful for various cross-spectrum techniques, as well as for jack-knives and null tests.

The mapmaker then uses the accept list and split list (if splits are requested), provided by *accept mod*, and creates a set of two maps of the accepted data for the each of the splits, as well as the map of all the accepted data.

3.7. Mapmaking

The mapmaker `tod2comap` produces temperature sky maps for each frequency channel. It takes in the processed TOD and the pointing information. The TOD can be written as a sum of the signal \mathbf{s} and the noise \mathbf{n} ,

$$\mathbf{d} = \mathbf{P}\mathbf{s} + \mathbf{n}, \quad (27)$$

where \mathbf{P} is the pointing matrix, which connects each time sample to a pixel on the sky. As it is the signal itself we are interested in, we have to rewrite this equation. Assuming that the noise is Gaussian, the log-likelihood function of the noise can be written as

$$\log \mathcal{L} \propto (\mathbf{d} - \mathbf{P}\mathbf{s})^T \mathbf{N}^{-1} (\mathbf{d} - \mathbf{P}\mathbf{s}) \quad (28)$$

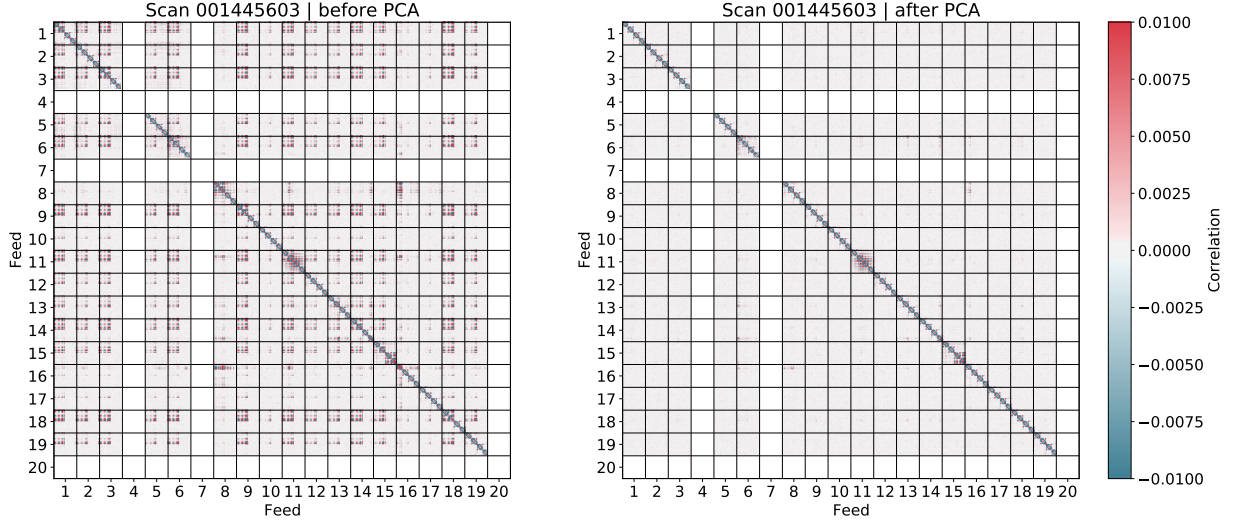


Figure 13: Correlation plot between each frequency in each feed before (left) and after (right) the filtering. It is evident that the filters are able to remove a lot of the correlated noise.

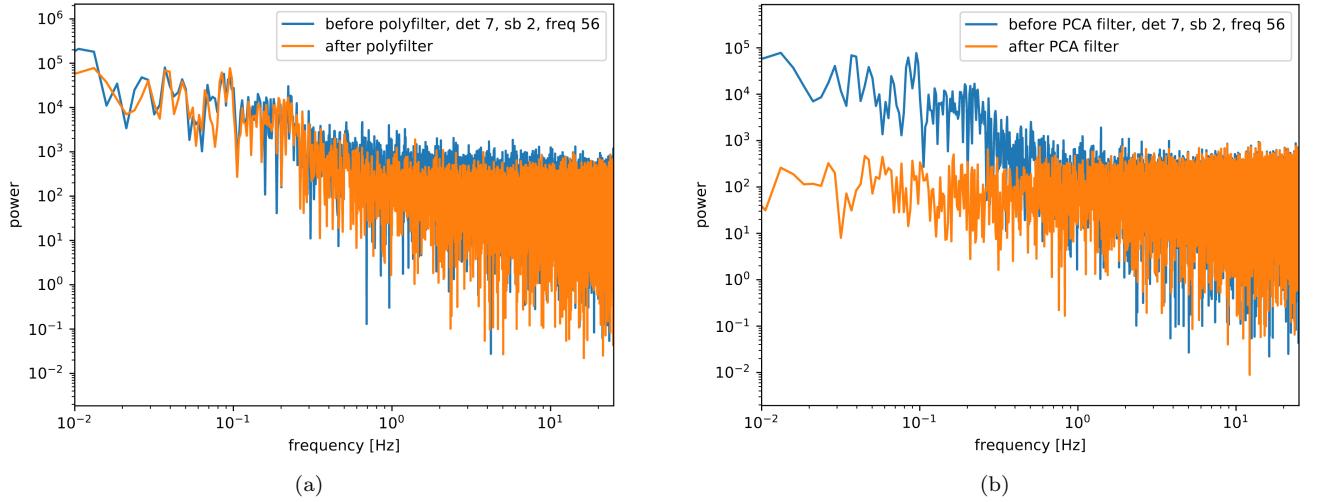


Figure 14: How the noise power spectrum changes during the processing. In both figures, the blue curve is before and the orange is after (a) the polyfilter and (b) the PCA filter. The noise goes from $1/f$ noise to white noise.

where \mathbf{N} is the covariance matrix. Setting the derivative of this equation to 0, we obtain the standard mapmaker equation

$$\hat{\mathbf{s}} = (\mathbf{P}^T \mathbf{N}^{-1} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{N}^{-1} \mathbf{d}. \quad (29)$$

As the data filtering explained above removes most of the correlated noise, we assume that the noise is white, making \mathbf{N} diagonal. We can then sum over the noise weighted data,

$$\hat{s} = \frac{\sum_i \text{data}/\sigma_i^2}{\sum_i 1/\sigma_i^2}. \quad (30)$$

Here σ_i is the noise of the sample, thus the samples with lower noise are weighted higher than the samples with

higher noise. The corresponding noise vector is then

$$\mathbf{n} = \frac{1}{\sqrt{\sum_i 1/\sigma_i^2}}. \quad (31)$$

As more data is collected, it might be necessary to introduce a more complex mapmaker. One possibility is to employ the conjugate gradient (CG) method.

4. ERROR PROPAGATION, SIMULATIONS, AND NULL-TESTS

The polyfilter induces some anti-correlations within each sideband...

4.1. White Noise Simulations

[Jowita] In order to assess the accuracy of the signal detection across cosmological scales, it is essential to estimate the uncertainty of experimental observables. We consider the spherically averaged power spectrum of the CO signal,

$$P(k) = \frac{\langle |f_{\mathbf{k}}|^2 \rangle}{n_x n_y n_z} V_{\text{vox}}, \quad (32)$$

extracted from the temperature sky maps by, first, taking the 3D Fourier transform of the (pixel-ordered) data. The absolute squared values of the Fourier coefficients are thereafter binned according to the magnitude of corresponding wave number, k , and averaged over all the contributions to a given k -bin. Finally, they are multiplied by the co-moving voxel volume, V_{vox} , and divided by the total number of voxels in each direction, $n_x n_y n_z$.

The main source of uncertainty in the power spectrum is the uncorrelated thermal noise with the standard deviation (noise per voxel) given by the radiometer equation (eq. 4/13), σ_N . Based on this value, determined in the previous stages of the pipeline, we can perform series of (typically $n_{\text{sim}} = 10$) white noise simulations to estimate error propagation. Each such simulation consists of creating a white noise map of samples drawn from the Gaussian distribution with standard deviation σ_N . This map is subsequently weighted by σ_N^{-2} (normalized with the mean), and its 3D power spectrum is calculated. The resulting dataset contains then n_{sim} values of the noise power spectrum for each k -bin. Computing the standard deviations across all simulations gives error bars in the final power spectrum extracted from the temperature map.

Moreover, the mean power spectrum of N temperature maps can be calculated as

$$\bar{P}(k) = \frac{\sum_i^N \sigma_{i,k}^{-2} P_i(k)}{\sum_i^N \sigma_{i,k}^{-2}}, \quad (33)$$

where $P_i(k)$ and $\sigma_{i,k}$ are the spherically averaged power spectrum of a single map and its uncertainty, respectively. The associated error bars at a given scale, k , become

$$\bar{\sigma}_k = \left(\sum_{i=1}^N \frac{1}{\sigma_{i,k}^2} \right)^{-1/2}. \quad (34)$$

[I should probably mention distinction between auto PS and xs at some point - but maybe this is subsection 4.3 (?)]

4.2. Transfer Functions

[Nils]

In order to estimate the efficiency of the analysis pipeline on different scales \vec{k} one can estimate the so-called filter transfer function. Because uncertainties in

the final power spectra of a map are not only dependent on the statistical error, but also subject to the pipeline efficiency, one needs to estimate the transfer function to obtain correct uncertainty estimates. For instance if the pipeline has a low efficiency for a given scale \vec{k} we expect the error bars to be larger, and vice versa.

We can estimate this transfer function by using a simulated signal that is added to the existing TOD of the telescope, and compare it to the pipeline output. In order to generate such a new TOD P_{new} with added simulations from a data cube, we use the pointing of the telescope to add additional signal with brightness T_{sim} to the existing TOD P_{old} such that

$$P_{\text{new}} = P_{\text{old}} + P_{\text{sim}} = k_B G \Delta \nu T_{\text{sys}} \left(1 + \frac{T_{\text{sim}}}{T_{\text{sys}}} \right). \quad (35)$$

From the TOD with added simulations P_{new} we then generate a map using the pipeline, and estimate the transfer function by comparing the power spectrum of the pipeline output to that of the input simulation. The output (noise weighted pseudo-) power spectrum P_{out} as a function of scale $\vec{k} = (k_{\parallel}, k_{\perp})$ is modeled as

$$P_{\text{out}}(\vec{k}) = T(\vec{k}) P_{\text{sim}}(\vec{k}) + P_{\text{noise}}(\vec{k}), \quad (36)$$

where P_{sim} is the power spectrum of the simulation as seen by the telescope, T is the transfer function and P_{noise} is the noise power spectrum. From this one easily finds the transfer function to be

$$T(\vec{k}) = \frac{P_{\text{out}} - P_{\text{noise}}}{P_{\text{sim}}}. \quad (37)$$

However, the modelling of the noise power spectrum is non-trivial and is highly dependent on the noise properties of the final maps. For now we estimate P_{noise} as the power spectrum of the original TOD before adding simulated signal, since these are noise dominated and have the same noise properties and systematics as P_{out} .

The effect of the different filters in **12gen** on the transfer function can be seen in Fig. 15, which shows a series of 1D transfer functions for different levels of filtering. One can see that the default settings, i.e. a polyfilter of first order and PCA filtering turned on, yields almost the same transfer function as the case where the PCA filter is turned off. The PCA filter is not expected to remove much of the input actual signal, as it only removes the components of the TOD that are the most correlated over frequencies and feeds thus potentially removing only the largest structures of the input signal from the simulation. When it comes to the polyfilter there is, however, some differences seen from the default settings. Using a polyfilter of $\mathcal{O}(1)$ or $\mathcal{O}(1)$ a considerable fraction of the input signal is removed by the pipeline on

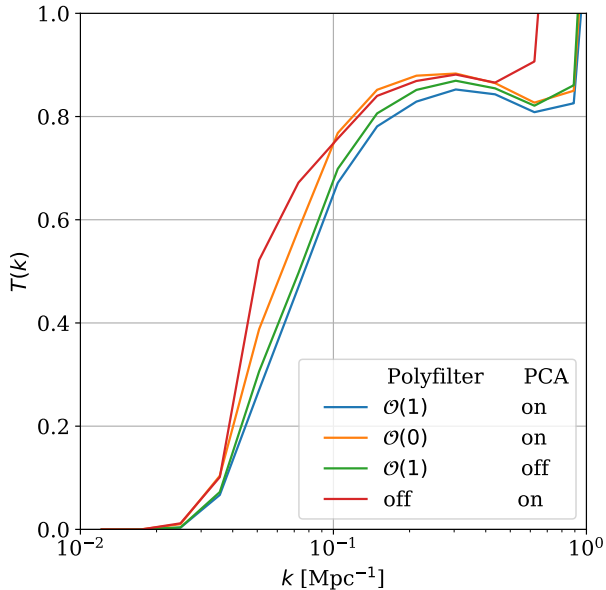


Figure 15: The 1D transfer functions $T(k)$ for different combinations of filtering in `12gen` as a function of scale k .

scales above $k \sim 4 \cdot 10^{-2} \text{ Mpc}^{-1}$. We see that the $\mathcal{O}(0)$ polyfilter yields a similar result as without the polyfilter near the peak regions, however a significant portion of signal in the $k \sim 4 \cdot 10^{-2} \text{ Mpc}^{-1}$ to $k \sim 1 \cdot 10^{-1} \text{ Mpc}^{-1}$ range is taken out when turning on the $\mathcal{O}(0)$ polyfiltering. The low transfer efficiency on low k for any of the shown filter combinations is likely due to the highpass filtering done in normalizing the TOD by its running mean.

When looking at the 2D version of the transfer function, as seen in Fig. 16, the effects of the polyfilter on the transfer function become more evident because the information in the spectral dimension ($_{\parallel}$) is not lost. As the polyfilter is designed to remove the $1/f$ -noise as well

as continuum foreground emission along the frequency dimension on each sideband, we expect the changes in the transfer function to be most visible in the large scale k_{\parallel} . This is indeed exactly what is seen in the difference between the transfer functions without and with $\mathcal{O}(1)$ polyfilter, $\Delta T(k)$ in Fig. 16 for low k_{\parallel} , where we note a 50 – 90% relative loss in power when using a $\mathcal{O}(1)$ polyfilter.

[If we haven't found out the reason for the bright spots at high k_{\perp} until the due date include the content of the following brackets] [The bright regions with $T(\vec{k}) > 1$ in the three plots of Fig. 16 at high k_{\perp} , corresponding to the sharp upwards turn of the graphs in Fig. 15 at high k , are not yet fully understood. Because the region in which this occurs is in the noise dominated regime, a potential cause could be that the noise power spectrum approximation we used here does not hold in this regime. However, more research is needed to fully understand this behaviour at high k_{\perp} fully.]

4.3. Null-tests

[Jowita/Håvard/Marie]

Test the quality of the data. Looking for sidelobes. Jackknives? Even vs odd numbered obsIDs, should not be a difference. Night data vs day data, only difference if day data are generally bad. Good vs bad feeds, feeds that generally performs well or bad, would expect some difference. Reveal possible problems with the data.

5. CONCLUSIONS

We have described the pipeline used to analyse the first 6 months of COMAP data (for the CO fields), going from raw data to temperature maps.

We thank ... Support from the Research Council of Norway through grant 251328.

REFERENCES

- Li, T. Y., Wechsler, R. H., Devaraj, K., et al. 2016, ApJ, 817, 169
- Weiland, J. L., Odegard, N., Hill, R. S., et al. 2011, ApJS, 192, 19
- Kovetz, E. D., Viero, M. P., Lidz, A., et al. 2017, arXiv:1709.09066
- Hill, G. J., Gebhardt, K., Komatsu, E., et al. 2008, Panoramic Views of Galaxy Formation and Evolution, 399, 115
- Papovich, C., Shipley, H. V., Mehrtens, N., et al. 2016, ApJS, 224, 28

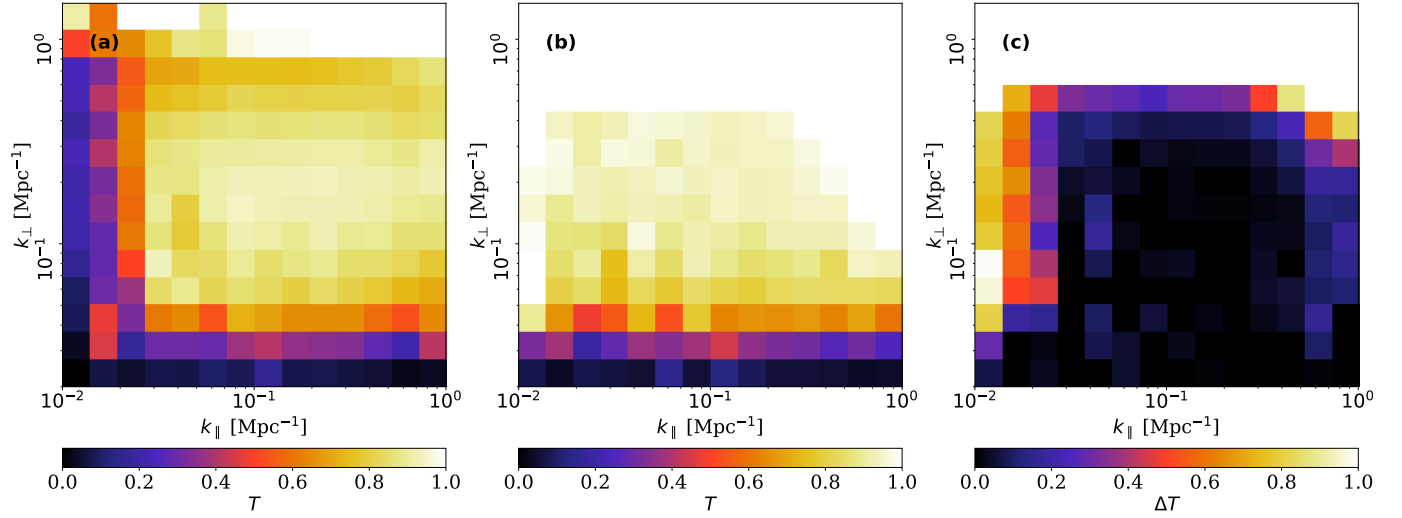


Figure 16: The 2D transfer function estimate of the default settings (a) in 12gen (i.e. with polyfilter of $\mathcal{O}(1)$ and PCA filter turned on), the transfer function estimate when not using the polyfilter (b) and the difference between the transfer functions without and with polyfilter turned on (c) all as function of scale $\vec{k} = (k_{\parallel}, k_{\perp})$.