Derivation of sideband gain ratio for *Herschel*/HIFI[★]

Do Kester¹, Ronan Higgins², and David Teyssier³

¹ SRON Netherlands Institute for Space Research, PO Box 800, 9700 AV Groningen, The Netherlands
e-mail: dokester@home.nl

² I. Physikalisches Institut der Universität zu Köln, Zùlpicher Straße 77, 50937 Köln, Germany

³ European Space Astronomy Centre, Camino Bajo del Castillo s/n, Urb. Villafranca del Castillo, 28692 Villanueva de la Cañada, Madrid, Spain

Received 19 August 2016 / Accepted 19 November 2016

ABSTRACT

Context. Heterodyne mixers are commonly used for high-resolution spectroscopy at radio telescopes. When used as a double sideband system, the accurate flux calibration of spectral lines acquired by those detectors is highly dependent on the system gains in the respective mixer sidebands via the so-called sideband gain ratio (SBR). As such, the SBR was one of the main contributors to the calibration uncertainty budget of the *Herschel*/HIFI instrument.

Aims. We want to determine the HIFI instrument sideband gain ratio for all bands on a fine frequency grid and within an accuracy of a few percent.

Methods. We introduce a novel technique involving in-orbit HIFI data that is bootstrapped onto standard methods involving laboratory data measurements of the SBR. We deconvolved the astronomical data to provide a proxy of the expected signal at every frequency channel, and extracted the sideband gain ratios from the residuals of that process.

Results. We determine the HIFI sideband gain ratio to an accuracy varying between 1 and 4%, with degraded accuracy in higher frequency ranges, and at places where the reliability of the technique is lower. These figures were incorporated into the HIFI data processing pipeline and improved the overall flux uncertainty of the legacy data from this instrument.

Conclusions. We demonstrate that a modified sideband deconvolution algorithm, using astronomical data in combination with gas cell measurements, can be used to generate an accurate and fine-granularity picture of the sideband gain ratio behaviour of a heterodyne receiver.

Key words. methods: data analysis – instrumentation: spectrographs – space vehicles: instruments

1. Introduction

The Heterodyne Instrument for the Far-Infrared (HIFI; [de Graauw et al. 2010](#)) was a high-resolution spectrometer flown aboard the *Herschel* Space Observatory ([Pilbratt et al. 2010](#)), which was operated by the European Space Agency between May 2009 and April 2013. The HIFI provided a complete spectral coverage over the 480–1272 GHz and 1430–1906 GHz ranges (625–240 and 208–157 μm , respectively) using a collection of heterodyne mixers offering instantaneous bandwidth of up to 4 GHz. Heterodyne detectors are in essence sensitive to two simultaneous portions of the electromagnetic spectrum called the upper and lower sidebands (USB and LSB). When no rejection (i.e. attenuation) of the signal from one of the sidebands is performed, the system is called double sideband (DSB), as was the case for HIFI. The net result of such a DSB mixing device is a spectrum, where the information from both the USB and LSB get folded together into one single spectrum at the down-converted intermediate frequency (IF). In the mixing process, the signal from either sideband is subject to separate instrument gains, which therefore impacts the absolute calibration of lines present in the respective sidebands differently.

The respective fraction of the DSB signal belonging to either of the sidebands is represented by the so-called sideband gain ratio (thereafter SBR; see also Sect. 2 for other conventions

used in this paper). The measurement of this instrument parameter is key for an accurate line flux calibration, yet it is particularly challenging. To first order, the broadband response of a heterodyne mixer can be assessed using Fourier transform spectroscopy (FTS; [Higgins et al. 2014](#)), whereby a monochromatic signal is radiated onto a mixer and the total power response of the mixer is recorded (e.g. [de Lange et al. 2003](#); [Teipen et al. 2005](#); [Karpov et al. 2007](#); [Cherednichenko et al. 2008](#)). How accurately the total power response measured by the FTS matches the heterodyne sideband ratio is a topic of debate, especially since FTS measurements are usually performed with an unpumped mixer; i.e. the mixer diode is not illuminated by any local oscillator signal and therefore not working at the same operating point as a pumped mixer. For example, [Higgins et al. \(2014\)](#) showed that some of the FTS measurements taken for various HIFI bands did not compare well with sideband ratio determined with a pumped mixer.

For pumped mixers, one possible approach is to temporarily transform the DSB into a single sideband (SSB) system by rejecting the signal from one of the sidebands, typically using a Martin-Pupplett interferometer, and measuring the response of each sideband independently (e.g. [Belloche et al. 2004¹](#)). Another technique consists of measuring the intensity of absorption lines through a laboratory gas cell. This method was successfully used to calibrate the SBR of the SWAS heterodyne mixers ([Tolls et al. 2004](#)), and was then intensively applied

[★] *Herschel* is an ESA space observatory with science instruments provided by European-led Principal Investigator consortia and with important participation from NASA.

¹ Available from <https://www.iram.fr/IRAMFR/PV/lowfreqs/report.ps.gz>

during the prelaunch validation campaign of the HIFI instrument (Teyssier et al. 2004; Higgins et al. 2010, 2014). The main drawback of this method, however, is that the SBR can only be assessed at the isolated frequencies where lines of a particular species emit in the gas cell. This is a noticeable restriction when considering the large frequency coverage over which the HIFI instrument had to be calibrated. Furthermore, it provides accurate SBR estimates for saturated lines, but brings additional uncertainties when applied to non-saturated lines, since a dedicated line emission model is needed to estimate the line opacities.

In either case, the derived calibration parameters rely entirely on laboratory measurements and can therefore hardly be revisited once the instruments are in space. In this paper, we describe a novel method combining those prelaunch measurements with the highly redundant astronomical data obtained in orbit to derive the SBR with a much finer frequency granularity. The method was successfully applied to all seven HIFI mixer bands and its outcome was the basis for the final calibration tables in use for the HIFI data processing pipeline.

In Sect. 2 we explain the conventions used in the paper; we present in Sect. 3 the data that were used. The novel method is described in Sect. 4; we also discuss the failure modes and the validation of the method using simulated data. In Sect. 5 we discuss the various cleaning operations on the data that were needed to get our results. The latter are shown in Sect. 6. In Sect. 7 we compare our results to additional evidences of the SBR characteristic derived from other instrument signatures, and a yet largely unexploited dataset measured with the HIFI gas cell on a methanol gas sample (Teyssier et al. 2004; Higgins 2011). Finally, in Sect. 8 we present the SBR values that were adopted in the HIFI pipeline.

2. Conventions

The concept of sideband gain ratio can vary in different astronomical facilities making use of heterodyne instruments, therefore it is important to clarify the conventions we use throughout this paper. From an engineering point of view, this parameter is the explicit ratio of the mixer response gains applicable to the respective USB and LSB frequencies, which we call here γ_{usb} and γ_{lsb} . The gain ratio R therefore is written as

$$R = \frac{\gamma_{\text{usb}}}{\gamma_{\text{lsb}}}. \quad (1)$$

In the HIFI formalism, the sideband ratio considered for the data calibration equation is normalised by the sum of the respective USB and LSB gains (Ossenkopf 2003; Roelfsema et al. 2012) and is defined as

$$G_{\text{usb}} = \frac{\gamma_{\text{usb}}}{\gamma_{\text{usb}} + \gamma_{\text{lsb}}}. \quad (2)$$

The value G_{usb} corresponds to the fraction of the DSB signal belonging to the USB, while $G_{\text{lsb}} = 1 - G_{\text{usb}}$ is the fraction of the DSB signal belonging to the LSB (Higgins et al. 2014). The values G_{usb} and R are simply related by

$$G_{\text{usb}} = \frac{R}{1 + R}. \quad (3)$$

The method described in this paper consists in deriving the deviations of the respective gains compared to those of a perfectly gain-balanced system ($R = 1$, or $G_{\text{usb}} = 0.5$). In the remainder of the paper, we refer to this sideband gain deviation as delta gain,

(δ_g in formulae). In our formalism, the respective normalised USB and LSB response gains can be written as $(1 + \delta_g)$ and $(1 - \delta_g)$, and we have

$$G_{\text{usb}} = \frac{1 + \delta_g}{2}. \quad (4)$$

3. Data

The novel sideband ratio extraction technique is exercised against two sets of data. The first set stems from HIFI prelaunch test data, acquired with a laboratory gas cell (Higgins et al. 2014). The second set corresponds to spectral survey science data acquired by HIFI in orbit on bright submillimetre emitters (see Roelfsema et al. 2012 for a detailed description of the HIFI observing modes).

Both datasets contain independent data for the H and V polarisation. We treat the polarisations separately and derive sideband gains for H and V in all bands.

3.1. Laboratory data

Prior to launch, the HIFI instrument was extensively tested in the laboratory. Following the precepts of the SWAS prelaunch validation programme (Tolls et al. 2004), a dedicated gas cell measurement campaign was carried out to derive the sideband gain ratio over the HIFI frequency range (Teyssier et al. 2004). The main focus was put on simple molecules such as ^{12}CO , ^{13}CO , OCS or H_2O , which would saturate in the gas cell set-up. These species, however, provided lines only at a few discrete points in the HIFI frequency space, offering therefore an incomplete view of the overall SBR profile with frequency. Results from this campaign are summarised in Higgins et al. (2014). The big picture derived at the time was that of an essentially balanced sideband gain ratio in bands 3 to 7 (albeit with large error bars), while the gain ratio in bands 1 and 2 would experience a more pronounced deviation from a balanced system, especially at the band edges. Band 1 had peculiarities in itself with large in-band SBR variations (see also Sect. 6). In the present study, the simple molecule data were primarily used as priors to the delta gain derivation (Sect. 4.2).

In order to get an idea of the SBR behaviour on a finer frequency grid, a complete survey of a more complex and line rich gas sample (methanol, or CH_3OH) was performed at pressures between 0.15 and 0.25 mBar. Methanol data were taken over the entire frequency range of HIFI and account for 80% of the 6750 gas cell spectra taken. This dataset also provided representative HIFI data ahead of the *Herschel* launch and helped to validate the spectral deconvolution algorithm (Teyssier et al. 2005). Using the methanol lines to extract sideband gain ratio is more complex than for simple saturated lines, as a detailed radiative transfer analysis is required to accurately assess all individual line opacities. This is discussed in Sect. 7.2.

In addition to providing a test dataset for the delta gain determination, the methanol measurements themselves are of spectroscopic interest. They indeed represent the first intensity-calibrated spectra of methanol over such a wide frequency range in the submillimetre (the flux intensity was calibrated against reference black bodies). A full deconvolved methanol spectrum taken over all seven HIFI bands is provided as part of a dedicated ancillary product by the *Herschel* Science Centre².

² See <http://www.cosmos.esa.int/web/herschel/ancillary-data-products>

Table 1. Summary of usable data from in-orbit spectral scan observations for each considered astronomical object.

Source	Band							
	1		2		3		4	
	a	b	a	b	a	b	a	b
Orion KL	×	×	×	×	×	×	×	×
Orion S	×	×	×	×	×	×	×	×
Orion Bar	×	×	×	×	×	×	×	×
Sgr-B2(N)	×	×	×	×	×	×	×	×
Sgr-B2(M)	×	×	×	×	×	×	×	×
OMC-2	×	×	×	×	×	×	×	×
AFGL 2591	×	×	×	×	×	×	×	×
IRC+10216	×	×	×	×	×	×	×	×
LDN 1157			×					
iras 16293	×	×	×	×	×	×	×	×
NML Tau ^a	×	×	×	×	×	×	×	×
ngc 6334i	×	×	×	×	×	×	×	×
ngc 1333			×					
Vy CMa			×	×	×	×	×	×
II Lup			×	×	×	×	×	×
Mars			×				×	×
W75N							×	×
NGC 7027							×	×

Notes. ^(a) NML-Tau is a calibration observation.

3.2. In-orbit data

The in-orbit data used in this paper are a selection of spectral scans obtained towards the brightest sources observed by HIFI. Spectral scan observations consist in acquiring data at strictly irregular frequency spacing to build redundancy in the number of measurements featuring a given range of the sky frequency domain. This redundancy is key to resolve the DSB degeneracy and allows us to deconvolve the data to recover the SSB spectrum of the observed target; see also Sect. 4. The frequency spacings need to be irregular to avoid aliasing in the deconvolved spectrum (Comito & Schilke 2002).

Our selection criterion was to contemplate all observations covering the complete frequency range in a given HIFI band, and that feature lines with a signal-to-noise ratio of 3 and higher. For some objects, observations are available in all 14 HIFI bands, while others have a contribution in only one band. Table 1 lists all objects used here with names as found in the *Herschel* Science Archive (HSA).

4. Method

In order to determine how much the respective fluxes from the USB and LSB contribute to a measured DSB spectrum, we would need to know these fluxes to begin with. Since figuring out them is the intrinsic objective of any given astronomical measurement, those fluxes are generally unknown. We can, however, deconvolve a complete spectral scan observation with a balanced gain and use that as a proxy. By doing so, we interpret the residuals of the deconvolution problem as inputs for determining the delta gains. Consequentially we need very good data with little noise and bright lines where gain differences show most conspicuously. Noise here is defined as the standard deviation of the residuals left in the deconvolution process. It includes statistical (radiometric and drift) effects and systematic effects such as the delta gains, which we study in this paper, but also other effects (e.g. spurs) that we would like to reduce as much as possible (see Sect. 5) to make the delta gain effects more conspicuous.

4.1. Deconvolution

Deconvolution addresses the problem of constructing an SSB spectrum such that the sum of the folded USB and LSB flux contribution to any given frequency range equals that of the DSB spectrum. The spectrum, as a function of the frequency, $S(f)$, is represented as a table; each row has a flux value and a frequency value. Each flux value is a parameter to be estimated in the deconvolution process.

In an heterodyne instrument such as HIFI, two parts of the spectrum are mixed into the observed spectrum, and down-converted at the so-called intermediate frequency (IF) to form $F(\phi)$. The IF is defined as $\phi = |f - f_L|$, where f_L is the local oscillator (LO) frequency that the mixer is tuned to. Assuming a delta gain of size δ_g , we get

$$F(\phi) = (1 - \delta_g)S(f_L - \phi) + (1 + \delta_g)S(f_L + \phi). \quad (5)$$

With $\delta_g = 0$ the system is balanced in gain. In the deconvolution process we assume that δ_g is known.

The data, D , are taken at intermediate frequencies, generated by dozens of LO settings. Each data point D_i is directly comparable with the model point $F_i = F(\phi_i)$. We find the elements of the spectrum, $S(f)$, by minimizing χ^2 ,

$$\chi^2 = \sum_i (D_i - F_i)^2, \quad (6)$$

where the index i runs over all data points in a spectral scan observation. These observations were taken such that each part of the spectrum is seen about a dozen times, tapering to zero at the edges of the band. Only valid data (see Sect. 5) are considered in this process.

The problem in itself is linear in its parameters, i.e. all the spectral flux values. The sheer amount of pixels (10^5) per band prevents even simple matrix operations. We use the conjugate gradient fitter from the HIPE³ software (Ott 2011), which easily finds a minimum as the problem is still linear and thus monomodal. No maximum entropy is needed as the problem is well defined and well behaved.

One might wonder why Eq. (5) cannot be solved simultaneously for both the spectrum and delta gain. The answer is that this leads to unavoidable aliasing in the delta gain; the technical details are in explained in Sect. 4.4.2 and in Comito & Schilke (2002). We decided to stop after one cycle of deconvolution and delta gain derivations, as more cycles would increase the aliasing in the delta gain. Even with only one cycle, hints of aliasing can be seen in the results of bands 3 and above. See Sect. 6.

4.2. Prior knowledge

If we happen to have prior knowledge concerning what the delta gain might be, we can feed that into the deconvolution procedure. The USB and LSB contributions to the IF spectrum are multiplied by $(1 \pm \delta_g)$ before comparing them to the actual data. The resulting deconvolved spectrum then becomes a closer proxy of the true spectrum.

Our prior knowledge on the delta gains originates in the laboratory campaign described in Sect. 3.1. These priors are particularly needed when confronted with some of the drawbacks of our method, as described later in this section. Even if one only has a rough idea on what the (prior) delta gain might be, feeding

³ HIPE is a joint development by the *Herschel* Science Ground Segment Consortium, consisting of ESA, the NASA *Herschel* Science Center, and the HIFI, PACS, and SPIRE consortia.

it into the deconvolution and applying our method improves the delta gain and yields a result with much finer granularity.

4.3. Deriving the sideband gain

In principle the delta gains could be a freely varying function of each LO tuning and within the tuning it could be a freely varying function of IF. This complete problem is underdetermined and untractable as the number of unknowns is larger than the data size. For practical reasons we derive one value for every LO setting. We assume on theoretical grounds (Higgins et al. 2014) and based on the behaviour of some bright lines that δ_g depends linearly on the location in the IF. In the remainder of this section, we process an observation piecewise, one LO setting at a time.

We assume that we have a deconvolved spectrum that can serve as a proxy for the true sky spectrum. We now rewrite the deconvolution Eq. (5), where this time δ_g is the unknown parameter. We consider the spectral contributions in USB and LSB as a given,

$$F(\phi) = (1 + \phi' \delta_g)U + (1 - \phi' \delta_g)L, \quad (7)$$

where $F(\phi)$ is the spectrum on an IF scale, and L and U are the fluxes in the USB and LSB respectively. The parameter δ_g is the deviation from balanced gains and ϕ' is a normalised IF: $\phi' = \phi/\hat{\phi}$, where $\hat{\phi}$ is the midpoint of the IF range obtaining the linear dependency we need. Defining ϕ' this way, δ_g is also the average over the IF range. Assuming that the overall flux calibration is correct, the delta gain δ_g is added to the gain of one sideband and subtracted from the other.

With a normal distribution on the errors between the data, D_i , and the model, $F_i = F(\phi_i)$, we can write the likelihood as

$$P(D|\delta_g\sigma) = \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^N \exp\left(-0.5\frac{\chi^2}{\sigma^2}\right), \quad (8)$$

where P is the likelihood of the data given a value of δ_g and a value of the noise scale σ . The value χ^2 is the usual sum of the squared errors; cf. Eq. (6).

Equating the partial derivative of $\partial \log(P)/\partial \delta_g = 0$, after substitution of Eqs. (6) and (7), we obtain the maximum likelihood (ML) solution for δ_g

$$\delta_g = \frac{\sum D_i(U_i - L_i)\phi'_i + \sum L_i^2\phi'_i - \sum U_i^2\phi'_i}{\sum (U_i - L_i)^2\phi_i'^2}. \quad (9)$$

This equation has an interesting property. It is invariant under multiplication and addition to the data. When all data is scaled by a factor and offset by a term, the reconstructed spectrum also gets the factor and half of the term in each sideband. When D_i is replaced by $D'_i = aD_i + b$, the values of U_i and L_i resulting from the deconvolution process, are changed to $U'_i = aU_i + b/2$, resp. $L'_i = aL_i + b/2$. When we substitute these values in Eq. (9) both factor a and term b drop out. This invariance property reinforces our confidence in the derived delta gain in light of possible additive or multiplicative systematic errors in the overall calibration, for example from pointing errors.

Similarly we can derive the maximum likelihood value for the noise scale by differentiating the log of the likelihood of Eq. (8) to σ . It yields the noise variance

$$\sigma^2 = \frac{\chi^2}{N-1}. \quad (10)$$

The error in δ_g can be found from the noise variance divided by the inverse Hessian, which in this simple case is the denominator in Eq. (9). It results in

$$\sigma_\delta = \sqrt{\frac{\chi^2}{(N-1)\sum (U_i - L_i)^2\phi_i'^2}}. \quad (11)$$

We have now derived ML values for δ_g and for its standard deviation σ_δ for one IF spectrum. We can do this for all IF spectra in all available observations. As each IF is associated with a value of the LO, we can fit a cubic splines model to the delta gains as a function of LO. As the accuracy of the derived delta gain values is very different, we use a weighted fit where the weights equal the inverse of σ_δ . The results are the gains settings presented in Sect. 6.

4.4. Failure modes

The method assumes that the deconvolved spectrum can be used as a proxy for the real spectrum. Each point in the deconvolved spectrum is made from data that overlap it, in the LSB, and/or in the USB. The data in both the USB and LSB contain the contribution of the delta gains. It is assumed that these gain effects are averaged out in the deconvolved spectrum. There are two cases for which this assumption is not true.

4.4.1. Edges

At the edges of a deconvolved spectrum, the overlap only exists in the LSB (at the low end) or USB (at the high end). The averaging therefore only applies over one sideband and the gain effects are not averaged out.

In principle edge effects can be minimised by combining different (adjacent) bands into one spectrum. Spectral points at high frequencies in one band overlap with points at low frequencies in the next band. In practice additional issues appear because of potential calibration inconsistencies between observations, such as differences in optical coupling, random pointing errors, epoch-dependent contribution of reference position, etc. We should bear in mind that we are using residuals as our input data, so we are looking for second order effects. We investigated this option but finally abandoned it with the exception of the transition between band 1a to 1b.

4.4.2. Aliasing

Because the distance between the LSB and the USB is fixed at about 12 GHz, a periodic alias of about 24 GHz is possible in the delta gains.

If a delta gain is too high at some point in the LO, it results in a deconvolved flux that is too low at the USB frequency and too high at the LSB frequency; see Eq. (7). When the LO has moved by 12 GHz, the flux that was first seen in the USB is now seen in the LSB. Therefore a consistent situation emerges when the gain in the LSB is also too high, meaning that delta gain is too low. With this low delta gain we force the flux at the USB frequency to a value that is too high. This repeating process results in a periodic alias of 24 GHz in delta gain, which is completely consistent with the data.

This alias results in a periodically lower and higher deconvolved spectrum. As such any quasi-periodic function of 24 GHz will suffice. Our method is intrinsically blind for these features and any quasi-periodic delta gains with a period around 24 GHz should be carefully checked.

Fig. 1. Three successful reconstructions of analytic functions for the delta gain (upper 3 panels) followed by a failed reconstruction of a periodic gain using the alias period of 24 GHz. The reconstructed delta gain is shown in green, while the input delta gain is indicated in dashed blue. The lower panel shows the deconvolved simulated methanol spectrum in arbitrary units.

We attempted to solve Eq. (7) simultaneously for gains and fluxes, but this was not possible owing to this aliasing. An aliased gain appeared with a period of 24 GHz and an amplitude that is as large as possible. By keeping the gains fixed in the first round and the fluxes in the second round this problem was circumvented.

4.5. Validation on simulated data

Using the gas cell line fitting code from Higgins et al. (2010), together with entries from the JPL Molecular Spectroscopy Database (Pickett et al. 1998), it is possible to generate a synthetic methanol spectra with a known detector gain profile. This spectra can then be used to test the delta gain derivation method.

Using an actual observation as a template, we replaced the data with simulated data from the methanol spectrum at the proper LSB and USB frequencies and multiplied them with an analytically chosen delta gain function. We also added Gaussian noise of 1 K to the data. Given this dataset we constructed a deconvolved spectrum using a *balanced* gain. From the residuals of the deconvolution process, we extracted the delta gain according to the method described in Sect. 4. The resulting delta gain were then compared with the known input delta gain to assess the accuracy of the method.

We present four cases in Fig. 1, where we compare the reconstructed gain with the artificial gain folded into the simulated data. From the top down, we present four input delta gains: the first input has a constant positive delta gain in the upper panel;

the next input was sinusoidal, with arbitrary period; the third input is a broad parabolic gain; and, finally, in the 4th panel we introduce a periodic delta gain with the same period as the alias period. In the lower panel, we show the deconvolved simulated spectrum.

It is clear that the gain is well recovered in the upper three panels with the exception of some 12 GHz from the edges where the reconstruction starts to fail because of the edge effect, which is described in Sect. 4.4.1. Also, where there are fewer lines that are less bright, the reconstruction is less accurate. When we deliberately introduce a gain periodic with the alias period of 24 GHz, the gains cannot be reproduced anymore for the reasons explained in Sect. 4.4.2.

The outcome of the derivation exercise in the three first cases (and in several others not shown here) strengthened our confidence in the method.

5. Data cleaning

Because we are using the residuals of an earlier solution (i.e. deconvolution) as the data for the calculation of the sideband gains, it is of paramount importance that our data are as clean as possible. Our data were taken from the *Herschel* Science Archive (HSA)⁴ with the reprocessing version of HIPE 13.0. In this product version, the so-called electric standing waves (ESW) are removed from the data of bands 6 and 7 (Kester et al. 2014). All spectral scan data were inspected and flagged for spurs and bright lines. We eliminated the data flagged as spurs.

It was necessary to carry out several other correction and cleaning operations on the data before the deconvolution process and, subsequently, the gain derivation could be performed. These operations are described in the following.

5.1. Optical standing waves

The continua in our in-orbit data are generally too faint to see optical standing waves, which are proportional to the impinging flux, with the exception of the Mars observations (Table 1); in these observations, optical standing waves contribute to about 1% of the continuum and, most predominantly, at a period of ~100 MHz. In the laboratory methanol data, these waves contribute to a similar magnitude, but at a period of ~177 MHz in all bands. It is related to an optical element in the test set-up. Sinusoidal functions with this period were fitted to the data, masking lines, spurs and other blemishes and, finally, these functions were removed from the data.

5.2. Large-scale baseline correction

As each IF spectrum is largely overlapping with the next, alternating level jumps between subsequent measurements cannot be produced by the astronomical sources themselves. These jumps must be a type of remaining calibration error. The medians of subsequent IF spectra should follow some smooth global trend.

Disregarding both spurs and bright lines, medians were taken over all remaining data in each IF spectrum to estimate the continuum level. We modelled the medians as a cubic spline function in the LO value with one knot every 2 GHz. This leaves enough flexibility to follow the true continuum of the spectrum. The difference between the medians and spline model were

⁴ <http://www.cosmos.esa.int/web/herschel/science-archive>

subtracted from the IF data. The first deconvolution step was then performed from these corrected data (Sect. 4.1).

5.3. More spurs and other blemishes

Once we carried out the deconvolution, we summed the folded LSB and USB contributions and compared the resulting spectra with the data. We inspected the residuals carefully for outliers not associated with sideband gains. We still found more spurs, extra noisy regions and sometimes whole IF spectra that were unusable. All these cases were discarded and flagged as bad data. The deconvolution was then run again from this point. This process had to be iterated several times in some cases, as large spurs leave ghosts in the spectra, which caused outliers of their own.

5.4. Small-scale baseline correction

When all the bad data sections have been identified, there are still systematic deviations in the residuals. It appears that the residuals over one IF measurement curve in such a manner that cannot originate from the astronomical sources in the same sense as the previous baseline jumps. These systematic errors would essentially increase the error in the sideband gain derivation. We modelled these baseline curves with a four knot cubic spline model and subtracted them from the data. Then we started deconvolution again from that point.

5.5. Contaminated off-source positions

In the deconvolved spectra of Sgr-B2 (Table 1) we see absorption features that dip below zero. This points to the presence of some flux source in the reference position used to calibrated the data, which is not so surprising near the crowded Galactic centre. However, our method does not require strictly positive flux values and the algorithm also works on negative flux values. As the SBR equally affects the on-source and off-source flux, the results should be the same for the difference and this is not an issue in our problem.

6. Results

We applied the methods described in the previous sections to the datasets listed in Table 1. As (prior) gains in the deconvolution process we used a balanced gain, except in bands 1a, 1b, 2a, 5a, and 5b. The prior gains used in bands 1a, 1b, and 2a are shown as a green line in Fig. 2 together with delta gain measurement points derived from the laboratory data described in Sect. 3.1. The prior gains used in bands 5a and 5b are described in Sect. 6.5.

We present the results of the delta gain derivation in Figs. 3–5. Each panel represents the data for one band: H is on top and V is below. Sub-band a is on the left and b is on the right. In these figures each grey point represents a measurement of one LO setting at one given position on the sky. Blue points, present only in bands 1 and 2, originate from the methanol laboratory measurements. Darker points (both grey and blue) represent the derivation of the δ_g with a higher confidence level. The green dashed line shows the input (prior) delta gain. Finally, the red line shows the resulting delta gain, which is modelled as a cubic splines function with a (few) dozen equidistant knots. As more noise is present in higher bands, the number of knots decreases from about 35 knots in band 1 to about 12 knots in band 7.

Fig. 2. Prior gains used in the deconvolution of data from the bands 1a, 1b, and 2a. The delta gain derived from laboratory measurements of saturated lines of ^{13}CO , ^{12}CO , H_2O , and OCS are shown in orange crosses, red crosses, blue squares, and green circles, respectively. The prior gain over the applicable frequency ranges is shown as a green line.

At some frequencies there is a bright emission line present in almost all sky observations. Where such a line appears either in the LSB or USB, the delta gain can be derived with high accuracy; see for example the 570–580 GHz and 916–928 GHz ranges. At those locations the derived points are not only highly accurate, but they are also close together. As they originate from observations taken throughout the mission, it can be concluded that the delta gain did not change over time.

With increasing band number (i.e. higher frequencies) the figures show larger scatter in the derived data points. This is a consequence of two effects both working in the same direction: the data themselves are more noisy for higher bands and there are generally fewer and weaker lines in the spectra that could be used in those ranges.

We provide more details about the outcome of the delta gain derivation in each considered band in the following. In

Fig. 3. Delta gain derivation results for bands 1–3. Dots correspond to individual gain measurements; lighter-coloured dots have a lower level of confidence, and therefore a lesser weight in the fit. Black dots originate from in-orbit observations (Sect. 3.2), while mint crosses originate from laboratory methanol data (Sect. 3.1). The red line shows a smoothing cubic spline through the points. The green dashed line shows the prior gain. The black line represents a balanced mixer, i.e. delta gain = 0.

Sect. 7 we then discuss how these results match other independent evidences for particular delta gain profiles.

6.1. Band 1

The treatment of the band 1 mixer is the most complicated of all bands considered in the study. From an engineering point of view, the band 1 mixer is unique in HIFI as it uses a so-called twin-junction design, which in essence means that two different response gains are mixed into the global gain profile that

is applicable to the detector (Delorme et al. 2005). This implies that the gain structure is expected to be less shallow.

This behaviour was confirmed early on by the overall picture derived from the isolated saturated line measurement points (Fig. 2). The main manifestation of the perturbed response gain is in the ranges 560–580 GHz and above 600 GHz, respectively. In the lowest of those ranges, fine-granularity CO laboratory measurements have shown a continuous positive slope in the delta gain, which is corroborated by a highly accurate measurement point from H₂O around 564 GHz at significantly

Fig. 4. Delta gain results for bands 4–6. Dots are individual gain measurements; lighter dots are less precise and have less weight in the smoothing spline through the points. The spline is shown as a red line. The green line shows the prior gain. The black line represents a balanced mixer, i.e. delta gain = 0.

negative delta gain levels. The slope has to be inverted in order to match another highly accurate measurement point from OCS at 576 GHz, which indicates that the mixer gain ratio is balanced. The rest of the prior up to 580 GHz is then derived from the evolution of line intensities at different LO settings as inferred from spectral scan in-orbit data (Higgins et al. 2014). Above 600 GHz similar measurements from H₂O and OCS also argue for delta gain slopes variations, which we have turned into the prior profile shown in Fig. 2.

The lower end of band 1 (below 550 GHz) is better behaved and we make the assumption of a balanced gain all the way down

for the prior. We prefer a balanced prior when possible and here it made no difference whether we used a prior or not. The only exception is that of a positive delta gain prior that corresponds to an accurate gain ratio measurement in a water line at an LO tuning around 550 GHz.

Using the above prior, the derived delta gain largely follows the prior gains for frequencies above 550 GHz, but also reveals differences especially in areas where the prior was assumed to be balanced. This is particularly true in the range 500–530 GHz, where the delta gain experiences a plateau that is actually consistent with laboratory measurements of OCS around 500 GHz

Fig. 5. Delta gain results for band 7. Dots are individual gain measurements; lighter dots are less precise and have less weight in the smoothing spline through the points. The spline is shown as a red line. The green line shows the prior gain. The black line represents a balanced mixer, i.e. delta gain = 0.

(Fig. 2). Section 7.2 presents independent evidence that our results are a good representation of the delta gain behaviour in this mixer.

6.2. Band 2

It is known from broadband FTS measurements that the gain response of the band 2 mixer experiences a severe drop-off at the lower end of the tuning range (Higgins et al. 2014), making the LSB noticeably less efficient than the USB. In delta gain terms, this translates into a negative slope all the way from the lower end of the band to about 670 GHz. This profile is consistent with isolated delta gain laboratory measurements of OCS saturated lines. This slope was used as prior in the lower end of band 2, while the rest of the frequency range was assumed to be balanced.

Unsurprisingly, this shape is reproduced in the derived delta gain. If the prior were balanced in this lower edge, the delta gain would still show a departure from a flat response, however the absolute value of the gain imbalance would not be recovered and, therefore, would not match the laboratory measurement points from saturated lines; this is in line with what was deduced from the simulations (Sect. 4.5) and, in particular, that the method fails to recover delta gain structures at band edges. Above 670 GHz, the delta gain experiences a broad negative plateau up to ~710 GHz (less pronounced in V than in H). Above this point the delta gain appears essentially balanced.

6.3. Band 3

Band 3 uses a balanced prior gain all the way through the mixer tuning range. The derived delta gain come out essentially balanced as well, with the exception of a broad positive plateau between 810 and 845 GHz in the V polarisation and a narrower negative plateau between 910 and 950 GHz in the H polarisation.

6.4. Band 4

The approach for band 4 was similar to that of band 3 with a balanced prior throughout the whole frequency range. While the derived delta gain is essentially balanced in the V polarisation, until it reaches a narrow negative plateau above ~1080 GHz, the delta gain in the H polarisation experiences a regular drop from 0.02

at the lower end down to -0.02 at the upper end. We also note that parts of the data points suffered from significant noise in this band, which was mostly related to the frequency ranges in which the LO output power was deficient (see Sect. 7.1).

6.5. Band 5

Before we started this study, the delta gain of band 5 had already been derived based on the observation that the system noise temperature of this mixer featured a steady positive slope, which is indicative of a system that is constantly imbalanced (Higgins et al. 2014); see also the introduction of Sect. 7.1. The sideband gain ratio estimated by this method was $R = 0.94$ (Eq. (1)), which translates into $\delta_g = -0.03$.

We first derived the delta gain with our method using a balanced gain prior. The output of this computation was fully consistent with the picture of a broad plateau at -0.03 with the exception of the two edges, where the derived delta gain remained balanced. Knowing the limitations of our method of properly recovering the delta gain at band edges, we then fed in the constant negative plateau, at -0.03 , as prior and basically retrieved the very same profile over the complete band.

6.6. Bands 6 and 7

From engineering grounds (use of broadband antenna design and smaller IF bandwidth 2.4–4.8 GHz compared to 4–8 GHz for bands 1–5), it is expected that the mixers in use for bands 6 and 7, which are hot electron bolometers (HEB), exhibit a relatively flat response gain over their tuning ranges (Cherednichenko et al. 2008). Consequently we used a balanced gain as prior for both bands.

The resulting data are noisier than in any other band both because the detectors are less sensitive and less stable and owing to the scarcity of strong lines in the high excitation regimes required at these frequencies. Overall the derived delta gain is fairly balanced, which is in line with the expectation from the intrinsic gain response of those device. One could argue that band 6a shows evidence for a positive plateau in both polarisations, over most of its tuning range. Given the significant noise associated with these data and the lack of similar evidence in the other bands, we nevertheless interpreted the whole band 6 as a balanced system.

Fig. 6. System temperature vs. frequency for all 7 HIFI bands. The green lines show the H polarisation, while the blue dashed lines indicate the V polarisation. The a bands are on the left side of each panel and the b bands on the right. Grey areas indicate areas where the LO cannot pump the mixer at all, while pink areas show the LO-deficient areas.

7. Discussion

We compare here the results presented in the previous section to additional evidence for the delta gain behaviour observed by other methods or derived from ancillary data characterising the HIFI system.

7.1. System noise temperature

One of the main characteristics of an heterodyne radiometer is its sensitivity, which is measured in terms of so-called system noise temperature, or T_{sys} (e.g. Wilson et al. 2009) – the lower the T_{sys} the more sensitive the radiometer. Detailed T_{sys} measurements were performed in all phases of the mission, and their dependence with the LO tuning frequency is a direct indicator of the instrument performance. Figure 6 summarises the variation of T_{sys} with frequency for all seven HIFI mixer bands. When the LO output power is insufficient to pump the mixer, the instrument sensitivity is poor and the system noise temperature

becomes infinite; these areas of starved LO output power are indicated in grey in Fig. 6. A regime also exists in which the LO output power is just marginal, but the mixer is still sufficiently pumped to provide sensitivity to the detected signal. In these conditions the T_{sys} usually increases significantly, but is still measurable; these areas are shown in pink in Fig. 6.

At other places, where the LO provides sufficient output power for nominal pumping of the mixer, the T_{sys} variation with frequency can offer an additional hint concerning the SBR profile. Indeed, during the T_{sys} measurement, which consists of detector count measurements of two internal black-body sources, the LSB and USB contributions to the spectrum are measured with a potentially different T_{sys} and thus with different sensitivity (gain). When the T_{sys} is decreasing towards higher frequencies, the gain in the LSB is lower than that in the USB ($\gamma > 0$) and vice versa. What this means is that a constant, non-balanced gain, manifests as a constant slope in the T_{sys} versus LO frequency curves. As it turns out, we have also shown in Sect. 4.5

that the gain derivation method described in this paper is particularly accurate in identifying constant delta gains (except at the edges).

Although the T_{sys} profiles are not entirely determined by the sideband gain response, these profiles still provide an efficient proxy that can be used to reverse engineer how the delta gain is varying with frequency. The relationship between a constant T_{sys} slope and a flat non-balanced gain was first brought up in band 5 (Higgins et al. 2014). In this particular case, the value of the slope (6%) was directly linked to the actual gain ratio ($R = 0.94$). We consequently considered this band as a prototypical case of such an evidence and we looked for similar coincidences in the other bands. We show in the following how the various T_{sys} curves over the HIFI bands could be matched to the independent evidence of delta gain plateaus inferred from results given in Sect. 6.

7.1.1. Band 1

The band 1 delta gain plateau between 500 and 520 GHz matches well the negative T_{sys} slope observed in both polarisations in this range. Above 520 GHz, the T_{sys} rises again, however there is no measurable effect on the derived delta gain at least within the uncertainties. Above this range, the delta gain profile is too complex to be related to the simple behaviour of the system sensitivity.

7.1.2. Band 2

Both polarisation mixers of band 2 exhibit a negative delta gain plateau between 670 and 710 GHz. This range coincides with a positive T_{sys} slope, which is actually more pronounced in the H than in the V polarisation and is in agreement with the delta gain plateau level derived from our method. Above this, the T_{sys} is essentially flat, which is in line with the balanced delta gain picture we derived.

7.1.3. Band 3

If one disregards the isolated areas of deficient LO output power, the T_{sys} is essentially flat up to 910 GHz for the H polarisation, and then experiences an steady rise from that frequency onwards. In the V polarisation, in contrast, the T_{sys} exhibits a negative slope between 810 and ~840 GHz and essentially flattens further up. These ranges nicely match the respective areas of balanced and non-balanced constant delta gain plateaus derived in this band.

7.1.4. Band 4

The T_{sys} in the V polarisation can be assimilated as a flat curve in the first half of the band, and then exhibits a positive slope starting around 1070 GHz. This feature coincides well with the location of the negative delta gain plateau observed in this polarisation above 1080 GHz. In the H polarisation, on the other hand, it is hard to define clear segments of particular slopes, instead the overall shape resembles more that of a negative parabola. As it turns out, the cubic spline fit to the derived delta gain points in this case shows a behaviour that steadily decreases over the whole band. The analogy between linear T_{sys} slope and constant delta gain plateau would hold if one considers that the T_{sys} represents to first order the derivative of the delta gain response. As

such, we interpret the parabolic shape of the T_{sys} as consistent evidence for a constant negative slope in the delta gain profile.

7.1.5. Bands 5, 6, and 7

Band 5 was our prototypical case behind the relationship between constant sideband gain and T_{sys} slope. Consequently, their match in the HIFI data is obtained by construction.

In bands 6 and 7, the system noise temperatures are scattered, essentially because of the difficulty in optimally pumping the HEB mixers at all tuned frequencies. The big picture is that of a relatively flat sensitivity curve in both polarisations. One could argue, however, that the overall T_{sys} envelope in band 6a exhibits a negative slope in the first 50 GHz. Interestingly, we discussed in Sect. 6.6 the possible hint for a positive delta gain plateau in this band. It is hard to say whether these two pieces of evidence are related, as there is significant noise associated with the data used in the delta gain derivation for those bands.

7.2. Modelling of gas cell methanol data

The gas cell methanol measurement described in Sect. 3.1 can also be used to derive sideband gain ratio, albeit with a larger uncertainty compared with those based on saturated lines. While for the latter the SBR can directly be extracted from the line peak, the analyses of methanol data require estimates of the line opacity, since only a handful of those lines saturate at the HIFI frequencies. The line opacity is calculated based on the line strength (taken from the JPL catalogue), gas cell path length, gas temperature and pressure, line frequency, and line pressure broadening parameter; see Higgins (2011). The pressure broadening parameter is, however, not available for every transition of methanol. With dedicated observations this parameter can be determined by measuring spectra at various pressures; see Giesen et al. (1992) for a detailed discussion. Nevertheless this level of spectroscopic testing was beyond the scope of gas cell test measurements of HIFI and would have substantially increased the amount of test time required. Furthermore, the sideband ratio and pressure broadening are degenerated and both affect the line profile in a similar fashion: increasing the pressure broadening parameter behaves similarly to an increasing sideband ratio.

To circumvent this limitation, one can estimate the pressure broadening for transitions falling at frequencies close to those measured with one of the saturated gases (e.g. CO, H₂O, and OCS) for which the SBR has been accurately determined; see Higgins et al. (2014). The sample pressure is one key parameter to use in combination with this. The pressure gauges used in the test set-up did not provide a pressure measure that was calibrated absolutely out of the box, and the sensor zero level needed to be calibrated prior to any measurement campaign. This calibration was unfortunately not performed for every phase of the test campaign, leaving the calibration offset unknown for a fraction of the laboratory data and, in particular, at low pressures measurement such as methanol. We could recover this unknown by rescaling the pressure by an offset that is adjusted in such a way that it would reconcile the delta gain derivation from data taken at redundant frequencies over several epochs of the measurement campaign; it would also reconcile this derivation for saturated species for which the delta gain are known independently from the gas pressure.

The delta gain derived with this method are shown in Fig. 7. The relatively large error bars essentially stem from the additional uncertainty associated with the a posteriori

Fig. 7. Delta gain extracted from CH₃OH gas cell observations (black crosses and error bars) for the H and V polarisation detectors in bands 1 and 2 compared with the delta gain derived from the deconvolution method (red lines).

recalibration of the pressure sensor zero level. The delta gain derived from the novel method introduced in this paper are overlaid as a red curve. The agreement between the two is very good and confirms the assumptions made on the priors for bands 1 and 2 at a fine frequency granularity level. The most noticeable discrepancy is seen around 540 GHz in band 1a, where the methanol data suggest a delta gain dip that cannot be reproduced by the deconvolution method. Error bars are larger in this area and isolated measurements from saturated lines (Fig. 2a) do not fully support such negative delta gain in this range. Consequently we decided to not reflect this local trend into our priors.

As stated in Sect. 3.1, laboratory measurements of methanol were taken over the entire HIFI data frequency range. Similar to bands 1 and 2, A detailed radiative transfer approach to extract the side band ratio from bands 3–7 was undertaken. For the diplexer bands (bands 3, 4, 6 and 7) the degree of scatter in the fitted side band ratios were too high to provide any useful comparison to the deconvolution extracted sideband ratio. This additional scatter is associated with the diplexer mistune problem seen during the ground-testing campaign; see Higgins et al. (2014) for more information.

8. Conclusions

The delta gain derivation presented in this study was used to draw the final SBR calibration tables applying to the data processing pipeline of the HIFI instrument. We use an interpolation scheme because the SBR needs to be applied at any frequency tuned by the instrument. This implies that the SBR curves have to be relatively smooth and that these curves at least avoid

discontinuity on very short scales. To this aim, we took the cubic spline fits to the data points shown in Figs. 3–5 as a started point. The other requirement is that the tabulated SBR need to be realistic. These ratios indeed represent the characteristic of a physical device where certain design constraints cannot be compatible with a SBR behaviour that is too exotic. In particular one has to bear in mind that, although datasets were analysed by separating the a and b frequency regions of a given mixer band, they both belong to the same detector with its overall gain response function. Finally, the calibration tables need to come with an uncertainty. This is particularly relevant for HIFI, as the absolute flux calibration uncertainty estimate is based on an error propagation equation, where the SBR can contribute a sizeable percentage of the total uncertainty (Ossenkopf 2015).

In Fig. 8 we present the final sideband ratio tables adopted for the HIFI flux calibration (black lines), together with the assumed 1σ uncertainties on those numbers (dashed lines). The cubic spline fits originating from Figs. 3–5 are also shown (red lines), while the pink envelope represent the 1σ confidence levels applying to those SBRs; these confidence levels were not shown in the previous figures to avoid cluttering. See the caption for further details on the band organisation per panel.

Our guidelines in arriving at the final sideband ratios were that we adopted the prolonged stretches of deviations from a balanced gain, especially when they were confirmed by additional pieces of evidence, as described in Sects. 7.1 and 7.2. Figures 3 and 4 occasionally show periodic structures in the delta gain, several of which are close to the expected aliasing period of 24 GHz described in Sect. 4.4.2. We interpreted those structures as unreal and they were not taken into account (i.e. flattened) in the final adopted sideband ratios. The confidence levels, however, were degraded accordingly to reflect the amplitude of those fluctuations, and therefore the higher uncertainty on the SBRs in such areas. Similarly, in areas where prior gains were used for the deconvolution process, the SBR error was adapted in order to reflect our estimated uncertainty on the priors.

The final picture is that of sideband gain ratios following relatively broad plateaus, which are linked by smooth and narrow segments. The uncertainty on those ratios varies typically between 1 and 4% depending on the band and frequency. In bands 6 and 7, the limited accuracy of the data available did not provide enough evidence for a convincing departure from the balanced sideband gains expected from theoretical grounds. As such we decided to stick to this picture, albeit with an increase of the sideband ratio uncertainty to reflect the dispersion of the derived SBR. A special mentioning is warranted for band 1. In this band, the adopted sideband ratio is a complex function that provides the best combined match to all the information we have at hand, namely the priors (in turn based on laboratory data), spectral scan data, and independent pieces of evidence from the T_{sys} measurements and methanol laboratory SBR derivation.

The novel method presented in this work is not specific to the HIFI data, and as such it could be used to extract sideband gain ratio from any other double sideband mixer in use or in development. One prerequisite to the approach, however, is that a minimum set of a priori knowledge about the mixer response gain should be available, especially when dealing with band edges. Laboratory gas cell measurements remain, in that respect, the most accurate complementary technique.

Acknowledgements. This work has benefited from numerous discussions within the HIFI Instrument Control Centre team. We would like to thank in particular J. Pearson, W. Jellema, and V. Ossenkopf-Okada for their important input at

Fig. 8. Sideband ratio G_{usb} (see Sect. 2) as a function of local oscillator frequency for all 7 HIFI bands. The red lines show the result of our fit with the 1σ confidence region shown as a pink envelope. The black lines represent the finally adopted values in the calibration tables. The dashed lines indicate the 1σ uncertainty. Each panel contains 2 sub-panels; the upper panel shows the H polarisation and the lower panel indicates V polarisation. In those panels we no longer distinguish between the a and b parts of the mixer band.

some stage of this study. We would also like to thank an anonymous referee for substantial and insightful comments that helped improve the final version.

References

- Belloche, A., Navarro, S., Sievers, A., et al. 2004, Test Observations at Frequencies Below 80 GHz at the IRAM 30 m Telescope, Memo, IRAM
- Cherednichenko, S., Drakinskiy, V., Berg, T., Khosropanah, P., & Kollberg, E. 2008, *Rev. Sci. Instr.*, **79**, 034501
- Comito, C., & Schilke, P. 2002, *A&A*, **395**, 357
- de Graauw, T., Helmich, F. P., Phillips, T. G., et al. 2010, *A&A*, **518**, L6
- de Lange, G., Jackson, B. D., Eggens, M., et al. 2003, in Proc. 14th Int. Symp. on Space Terahertz Technology, eds. C. Walker, & J. Payne, 447
- Delorme, Y., Salez, M., Lecomte, B., et al. 2005, in Proc. 16th Int. Symp. on Space Terahertz Technology, 444
- Giesen, T., Schieder, R., Winnewisser, G., & Yamada, K. M. 1992, *J. Mol. Spectr.*, **153**, 406
- Higgins, D. 2011, Ph.D. Thesis, Maynooth, National University of Ireland
- Higgins, R., Teyssier, D., Borys, C., et al. 2014, *Exp. Astron.*, **37**, 433
- Higgins, R. D., Teyssier, D., Pearson, J. C., Risacher, C., & Trappe, N. A. 2010, in Proc. 21st Int. Symp. on Space Terahertz Technology, 390
- Karpov, A., Miller, D., Rice, F., et al. 2007, *IEEE Trans. Applied Superconductivity*, **17**, 343
- Kester, D., Avruch, I., & Teyssier, D. 2014, in *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, AIP Conf. Proc., 1636, 62
- Ossenkopf, V. 2003, The HIFI intensity calibration framework, Tech. rep., ALMA Memo 442.1, University of Cologne
- Ossenkopf, V. 2015, The propagation of systematic uncertainties in the intensity calibration for HIFI, Tech. rep., University of Cologne
- Ott, S. 2011, in *Astronomical Data Analysis Software and Systems XX*, eds. I. N. Evans, A. Accomazzi, D. J. Mink, & A. H. Rots, *ASP Conf. Ser.*, **442**, 347
- Pickett, H. M., Poynter, R. L., Cohen, E. A., et al. 1998, *J. Quant. Spectr. Rad. Trans.*, **60**, 883
- Pilbratt, G. L., Riedinger, J. R., Passvogel, T., et al. 2010, *A&A*, **518**, L1
- Roelfsema, P. R., Helmich, F. P., Teyssier, D., et al. 2012, *A&A*, **537**, A17
- Teipen, R., Justen, M., Tils, T., et al. 2005, in Proc. 16th Symp. on Space Terahertz Technology, 199
- Teyssier, D., Dartois, E., Deboffe, D., et al. 2004, in Proc. 15th Int. Symp. of Space Terahertz Technology, Northampton, Massachusetts, 306
- Teyssier, D., Comito, C., Dartois, E., et al. 2005, *IAU Symp.*, **235**, 186
- Tolls, V., Melnick, G. J., Ashby, M. L. N., et al. 2004, *ApJS*, **152**, 137