HERSCHEL SPACE OBSERVATORY

HERSCHEL EXPLANATORY SUPPLEMENT VOLUME II

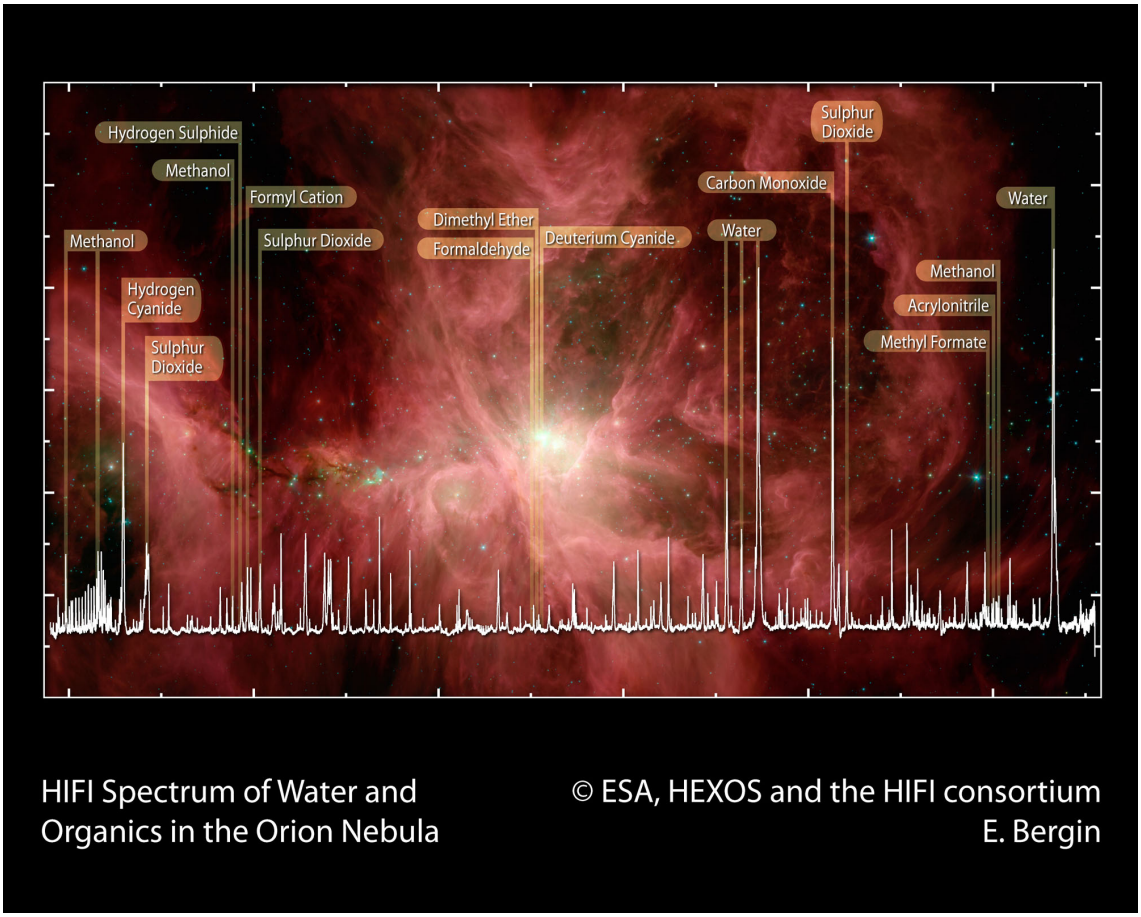

HIFI Spectrum of Water and Organics in the Orion Nebula

E. Bergin

THE HETERODYNE INSTRUMENT FOR THE FAR INFRARED (HIFI) HANDBOOK

HiFi ICC

HERSCHEL-HSC-DOC-2097, version 2.0, January 16, 2017

HIFI Handbook

Version 2.0, January 16, 2017

This document is based on inputs from the HIFI Consortium and the HIFI Instrument Control Centre.

Document editor and custodian:

Curator: David Teyssier, Herschel Science Centre, European Space Astronomy Centre, European Space Agency
Co-authors: Ian Avruch (SRON, The Netherlands), Sylvie Beaulieu (Univ. Waterloo, Canada), Jonathan Braine (LAB, France), Anthony Marston (ESA), Patrick Morris (NHSC, USA), Michael Olberg (OSO, Sweden), Miriam Rengel (ESA/ESAC and MPS, Germany), Russ Shipman (SRON, The Netherlands)

Contents

Chapter 1

Introduction

1.1 The Herschel Space Observatory

The *Herschel* Space Observatory (Pilbratt et al. 2010) is the fourth cornerstone mission in ESA's science programme. *Herschel* was successfully launched on 14 May 2009 from Kourou, French Guiana, and performed science and engineering observations until 29 Apr 2013 when the liquid Helium coolant boiled off. *Herschel* was in an extended orbit around the second Lagrangian point (L2) of the system Sun-Earth. After the end of operations, with the last command sent on 17 June 2013, the satellite was put into a safe disposal orbit around the Sun. *Herschel* telescope's passively cooled 3.5 m diameter primary mirror is currently the largest one ever flown in space. The three on-board instruments: the Heterodyne Instrument for Far Infrared (HIFI, de Graauw et al. 2010), the Photodetector Array Camera and Spectrometer (PACS, Poglitsch et al. 2010) and the Spectral and Photometric Imaging Receiver (SPIRE, Griffin et al. 2010) performed photometry and spectroscopy observations in the infrared and the far-infrared domains, from $\sim 60\,\mu$m to $\sim 672\,\mu$m. This spectral domain covers the cold and the dusty universe: from dust-enshrouded galaxies at cosmological distances down to scales of stellar formation, planetary system bodies and our own solar system objects.

A high-level description of the *Herschel* Space Observatory is given in Pilbratt et al. (2010); more details are given in the Herschel Mission & Satellite Overview Handbook. The first scientific results are presented in the special volume 518 of *Astronomy & Astrophysics* journal. Information with latest news, documentation, data processing and access to the *Herschel* Science Archive (HSA) is provided via the *Herschel* Science Centre web portal (`http://www.cosmos.esa.int/web/herschel/home`).

1.2 Purpose and Structure of Document

The purpose of this handbook is to provide relevant information about the HIFI instrument, in order to help astronomers understand and use the scientific observations performed with it.

The structure of the document is as follows: first we describe the HIFI instrument (Chapter 2), followed by the description of the observing modes (Chapter 3). The flux calibration scheme is explained in Chapter 4. The in-flight performance of HIFI is presented in Chapter 5. The high level HIFI data products available from the *Herschel* Science Archive are introduced in Chapter 6. The list of references is given in the last chapter.

1.3 Acknowledgements

This handbook is provided by the Herschel Science Centre, based on inputs by the HIFI Consortium, the HIFI Instrument Control Centre (ICC) and the NASA Herschel Science Center (NHSC).

1.4 Changes to Document

1.4.1 Changes in version 0.9

First major version prepared for internal distribution.

1.4.2 Changes in version 1.0

Implementation of feedback from members of the HIFI Instrument Control Centre.

Changed font size to 11 pt and re-organised figures to better match the corresponding text.

Updated Section 5.6 to account for band-dependent H/V co-alignment offsets.

Moved Section 4.4.4 to appendix A.

Updated Section 6.2 to reflected the latest situation with HPDPs and ADPs.

Updated Section 6.3.1 with latest numbers on science readiness statistics.

1.4.3 Changes in version 2.0

Implemented Section 5.5.

Added subsection 5.8.1.

1.5 List of Acronyms

ACA	Attitude Control Axis
ACS	Auto-Correlator Spectrometer
ACMS	Attitude Control and Measurement System
ADP	Ancillary Data Product
AGB	Asymptotic Giant Branch
AHF	Attitude History File
AME	Absolute Measurement Error
AOR	Astronomical Observation Request
AOS	Accousto-Optical Spectrometer
AOT	Astronomical Observation Template
APE	Absolute Pointing Error
CCD	Charge-Coupled Device
CSA	Calibration Source Assembly
COA	Common Optics Assembly
DBS	Double Beam Switching
DP	Data Processing
DRG	Data Reduction Guide
DSB	Double Sideband
EEF	Encircled Energy Fraction
ESA	European Space Agency
ESAC	European Space Astronomy Centre
ESW	Electrical Standing Wave
FCU	FPU Control Unit
FIR	Far Infrared Radiation
FOV	Field of View
FPC	Focal Plane Chopper
FPU	Focal Plane Unit
FSW	Frequency Switching
FWHM	Full Width at Half Maximum
GT	Guaranteed Time
GUI	Graphical User Interface
HCSS	Herschel Common Software System
HEB	Hot Electron Bolometer
HIFI	Heterodyne Instrument for the Far Infrared
HIPE	Herschel Interactive Processing Environment
HPBW	Half Power Beam Width
HPDP	Highly-Processed Data Product
HRS	High-Resolution Spectrometer
HSA	The Herschel Science Archive
HSC	The Herschel Science Centre (based in ESAC, ESA, Spain)
HSpot	Herschel Observation Planning Tool
HTP	HIFI Timeline Product
HUI	HSA User Interface
IA	Interactive Analysis

ICC	Instrument Control Centre
ICU	Instrument Control Unit
IF	Intermediate Frequency
ILT	Instrument Level Test (i.e. instrument ground tests w/o spacecraft)
IRTF	InfraRed Telescope Facility
ISM	Inter Stellar Medium
KP	Key Programme
LO	Local Oscillator
LOU	Local Oscillator Unit
LSB	Lower Sideband
LSR	Local Standard of Rest
LSU	Local Oscillator Source Unit
MO	Master Oscillator
MSA	Mixer Sub-Assembly
NHSC	NASA Herschel Science Center
OCXO	Oven Controlled Crystal Oscillator
OD	Observational Day
OSW	Optical Standing Wave
OT	Open Time
OTF	On-The-Fly
PACS	Photodetector Array Camera and Spectrometer
PSW	Position Switching
PV	Performance and Validation
RF	Radio Frequency
RMS,*rms*	Root Mean Square
RPE	Relative Pointing Error
SED	Spectral Energy Distribution
SAA	Solar Aspect Angle
SIAM	Spacecraft Internal Alignment Matrix
SIS	Semiconductor-Insulator-Semiconductor
S/N	Signal-to-Noise Ratio
SPG	Standard Product Generation
SPIRE	Spectral and Photometric Imaging REceiver
SRPE	Spatial Relative Pointing Error
SSB	Single Sideband
UKIRT	United Kingdom InfraRed Telescope
UPDP	User-Provided Data Product
USB	Upper Sideband
WBS	Wide-Band Spectrometer
WCS	World Coordinate System
YSO	Young Stellar Object

Chapter 2

The HIFI Instrument

2.1 Instrument Overview

2.1.1 What was HIFI

HIFI was the Heterodyne Instrument for the Far Infrared. It was designed to provide spectroscopy at high to very high frequency resolution over a frequency range of approximately 480–1272 and 1430–1906 GHz (625–240 and 208–157 µm, respectively). This frequency range was covered by 7 mixer bands, with dual horizontal and vertical polarisations, which could be used one pair at a time (see Table 2.1 for a detailed specification).

The mixers act as detectors that feed either, or both, of the two spectrometers connected to HIFI. An instantaneous frequency coverage of 2.4 GHz was provided for the high frequency band 6 and 7 mixers, while for bands 1 to 5 a frequency range of 4 GHz was covered. The data were obtained as Double Sideband (DSB) data, which means that each channel of the spectrometers reacts to two frequencies (separated by 4.8 to 16 GHz) of radiation at the same time (see Section 2.1.2). For many situations, this overlapping of frequencies is not a major problem and science signals are clearly distinguishable. However, particularly for complex sources containing a high density of emission and/or absorption lines, this can lead to problems with data interpretation. Deconvolution is therefore necessary for the data to create Single Sideband (SSB) data. This is especially important for spectral scans covering large frequency ranges on sources with many lines.

There were four spectrometers connected to HIFI, two Wide-Band Acousto-Optical Spectrometers (WBS) and two High Resolution Autocorrelation Spectrometers (HRS). One of each spectrometer type was available for each polarisation. They could be used either individually or in parallel. The Wide-Band Spectrometers covered the full intermediate frequency bandwidth of 2.4 GHz in the highest frequency bands (bands 6 and 7) and 4 GHz in all other bands. The High Resolution Spectrometers had variable resolution with subbands sampling up to half the 2.4 or 4 GHz instantaneous frequency range. These subbands had the flexibility of being placed anywhere within the 2.4 or 4 GHz range.

Band	LO frequency range (GHz)	Mixer type	IF bandwidth (GHz)	Typical HPBW at band centre
1a	487.5–553.5	SIS	4.0	$40''.4$
1b	562.5–628.5	SIS	4.0	$35''.3$
2a	634.0–718.0	SIS	4.0	$31''.1$
2b	722.0–794.0	SIS	4.0	$27''.7$
3a	807.0–852.0	SIS	4.0	$25''.3$
3b	866.0–953.0	SIS	4.0	$23''.1$
4a	957.0–1,053.0	SIS	4.0	$20''.9$
4b	1,054.5–1,114.0	SIS	4.0	$19''.4$
5a	1,116.2–1,236.0	SIS	4.0	$17''.9$
5b	1,235.0–1,272.0	SIS	4.0	$16''.8$
6a	1,430.0–1,558.0	HEB	2.4	$14''.1$
6b	1,578.0–1,698.0	HEB	2.4	$12''.8$
7a	1,701.0–1,794.0	HEB	2.4	$12''.0$
7b	1,793.0–1,902.0	HEB	2.4	$11''.4$

Table 2.1: Frequency tuning ranges of HIFI subbands.

2.1.2 How did HIFI work

Heterodyne detection

Sub-mm continuum radiation is best detected with bolometers, which act like thermometers, measuring the heat coming in and translating it to integrated intensities. Line radiation is much more difficult to detect. There are no amplifiers available to amplify the weak sky signals at sub-millimetre wavelengths. For lower frequencies there are, however, good amplifiers available, which can be small, low in energy consumption and weight. These are thus very suitable for a space observatory.

The solution is thus to bring the signal down in frequency, without losing its information content. This is accomplished through heterodyne techniques in which the sky signal is mixed with another signal (Local Oscillator (LO)) very close to the frequency of interest. In performing such mixing of signals, the resulting signal is converted to much lower frequency, while still having all the spectral detail of the original sky-signal. Modern mixing devices such as SIS (semiconductor-insulator-semiconductor) mixers or HEB (hot electron bolometer) mixers, not only perform the mixing but can also amplify the signal, making them eminently suitable for instruments like HIFI.

Mixing: The mixers used by HIFI are at superconducting temperatures (the HEBs are on the border of normal and superconducting), typically around 4 K. They are non-linear devices in that the output current is not directly proportional to the voltage across them – in fact their current-voltage curves have similarities to those of diodes. This allows amplification of the mixed signals of the incoming radiation and the on-board local oscillator (LO). These mixed signals are detected and amplified at the "beat" frequency ($\|\nu_{\rm sky} - \nu_{\rm LO}\|$) between each of the incoming source frequencies, $\nu_{\rm sky}$, and the single Local Oscillator frequency, $\nu_{\rm LO}$.

Intermediate Frequency (IF): The "beat" frequencies produce the so-called Intermediate Frequency (IF) of the instrument. Further amplification is made of these intermediate frequencies and, for HIFI, filtering allows the detection of IFs of 4 to 8 GHz which is done in the HIFI spectrometers.

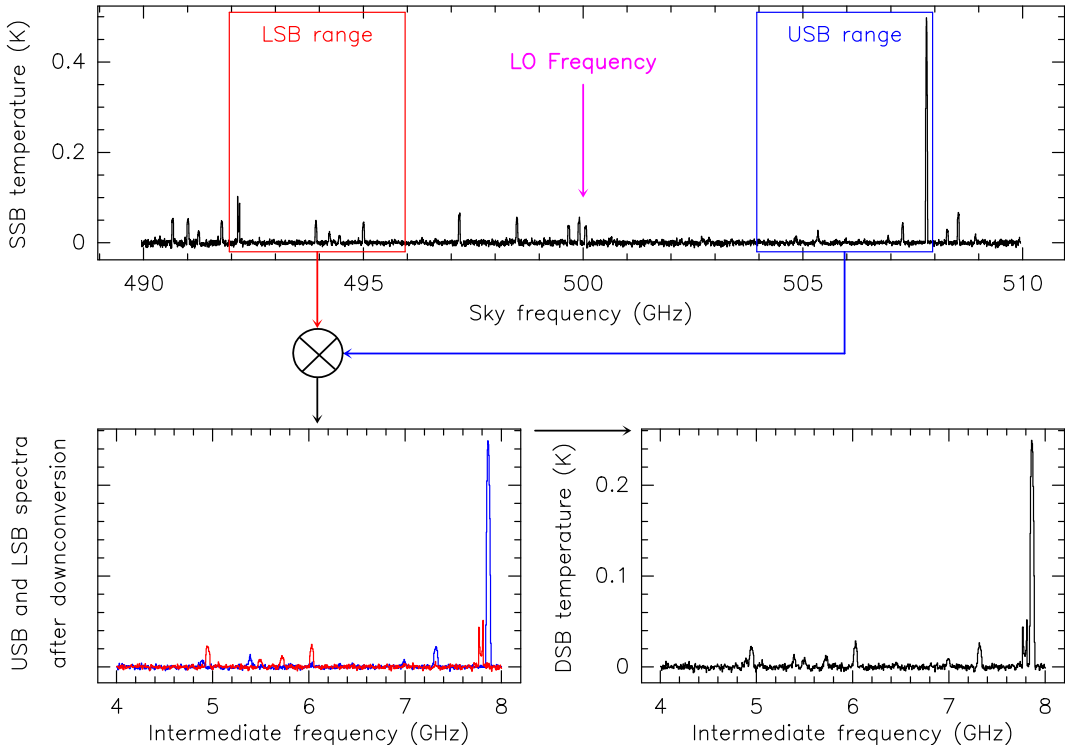

Figure 2.1: Illustration of the heterodyne mixing on an HIFI spectrum taken in band 1a on IRC+10216. The upper panel shows the SSB spectrum on a sky frequency scale (RF scale), together with the respective LSB and USB frequency ranges implied by an LO frequency tuning at 500 GHz, and an IF of 6 GHz. The lower left panel shows how the respective side band spectra (LSB in red, USB in blue) get combined in the down-converted spectrum at the intermediate frequency. Note how the LSB spectrum gets flipped around its frequency scale. The final DSB spectrum, shown in the lower right panel, is the sum of the respective USB and LSB spectra. Because the calibration scheme involves some normalisation to the emission of continuum black bodies, the respective line intensities will appear as half their SSB value (see Section 4.1.2).

Double Sideband data

In creating the intermediate frequency it should be noted that a given IF value (e.g., 5 GHz) can be obtained from a source frequency that is either 5 GHz higher or 5 GHz lower than the local oscillator frequency. If we consider this for a range of incoming frequencies we can see that the HIFI spectrometers measure two superimposed portions of an object's spectrum.

- The portion of the source spectrum 4 to 8 GHz (resp. 2.4 to 4.8 GHz for HEB mixers) above the LO frequency. This will be in ascending frequency order from $\nu_{\rm LO} + 4$ GHz to $\nu_{\rm LO} + 8$ GHz. This frequency range is called the *Upper Sideband* (USB).
- The portion of the source spectrum 4 to 8 GHz resp. 2.4 to 4.8 GHz for HEB mixers) below the LO frequency. This will be in descending frequency order from $\nu_{\rm LO} - 4$ GHz to $\nu_{\rm LO} - 8$ GHz. This frequency range is called the *Lower Sideband* (LSB).

This superposition is illustrated in Figure 2.1.

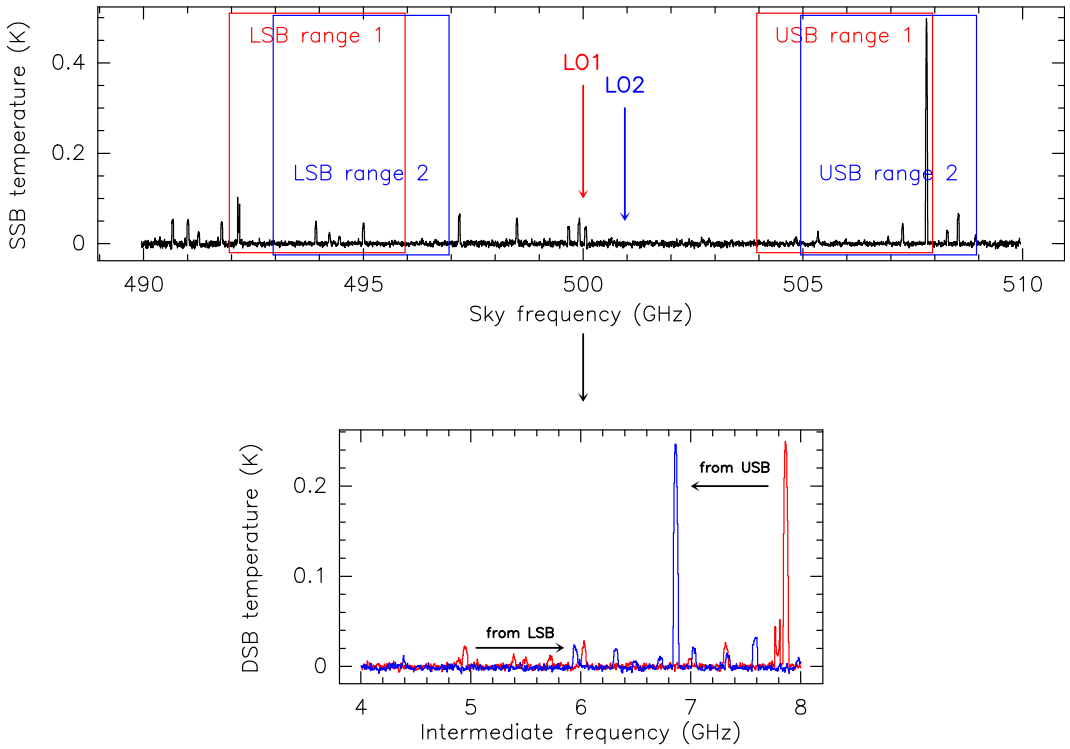

Figure 2.2: Same as Fig. 2.1, this time considering two nearby LO frequency tunings (LO1 at 500 GHz in red, LO2 at 501 GHz in blue). The superposition of the two separate DSB spectra in blue and red is shown in the lower panel. Note how the largest line from the USB goes down in the IF band frequency when the LO frequency is increased (compare with previous figure), while other weaker lines from the LSB go in the opposite direction.

For a number of regions where a single strong line of known frequency is the subject of study, knowing whether it is in the upper or lower sideband frequency range is easy to determine - and so it is easy to assign the correct frequency to the spectrum scale.

However, for cases where it is not known a priori which spectral lines are in which sideband the simplest way to determine this is by shifting the LO frequency. An increase in LO frequency will lead to USB features moving to lower IF frequencies and LSB features moving to higher IF frequencies (see Figure 2.2). It then becomes clear which sideband (and frequency) the features are in.

Deconvolution: Even the above technique becomes impossible for regions where there is a high density of spectral features. In such cases, the chances become quite high that USB features and LSB features will overlap and therefore blend. And the shifting of the LO may only lead to other feature overlaps. For this case spectral deconvolution techniques have been devised (see e.g. Comito & Schilke 2002). These allow large regions of frequency space to be sampled by changing the LO frequency in small steps. A reconstruction of the spectrum (called *Single Sideband*, SSB) can then be made.

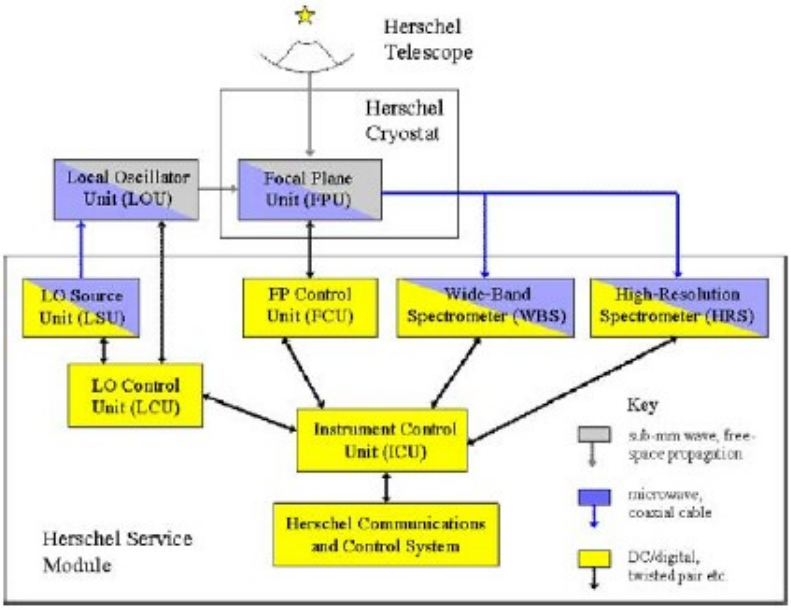

Figure 2.3: General HIFI component diagram.

2.2 Instrument Configuration

Referring to Figure 2.3, HIFI had five (hardware) sub-systems: the Local Oscillator and Focal Plane Sub-Systems; the Wide-Band and High-Resolution Spectrometers; and the Instrument Control Unit. Within the Local Oscillator Sub-System, a tuneable, spectrally pure 24–36 GHz signal was generated in the Local Oscillator Source Unit. This signal was then frequency-multiplied (upconverted) to 71–106 GHz, amplified, and further frequency-multiplied, by different factors for each of the LO chains, to the desired radio frequency (RF) in the Local Oscillator Unit. The result was a spectrally pure LO signal with a tuneable frequency and power level.

Fourteen multiplier chains covered the 480–1910 GHz range (625–157 µm), with two chains for each of the seven Focal Plane Unit mixer channels. The chain feeding the lower frequencies of the band was labeled ”a” whereas it was labeled ”b” for the higher frequencies (leading to the naming of mixer bands as 1a, 1b, 2a etc.).

The local oscillator beams were fed into the Focal Plane Unit through 7 windows in the Herschel cryostat. Within the Focal Plane Unit, the astronomical signal from the telescope was split into 7 beams. Each of these signal beams was combined with its corresponding LO beam, and then split into two linearly polarized beams that were focused into two mixer units. Each mixer unit generated an intermediate frequency (IF) signal that was amplified prior to leaving the Focal Plane Unit.

The IF output signals from the Focal Plane Unit could be coupled into two IF spectrometers: the Wide-Band Spectrometer, a four-channel (subband) acousto-optical spectrometer (AOS) that sampled the 4–8 GHz band at 1 MHz resolution; and the High-Resolution Spectrometer, a high-speed digital auto-correlator (ACS) that sampled narrower portions of the IF band at resolutions of up to 140 kHz.

Each of the spectrometers included a warm control electronics unit. These four control units were, in turn, commanded by a single Instrument Control Unit (ICU), which also interfaced with the satellite’s command and control system.

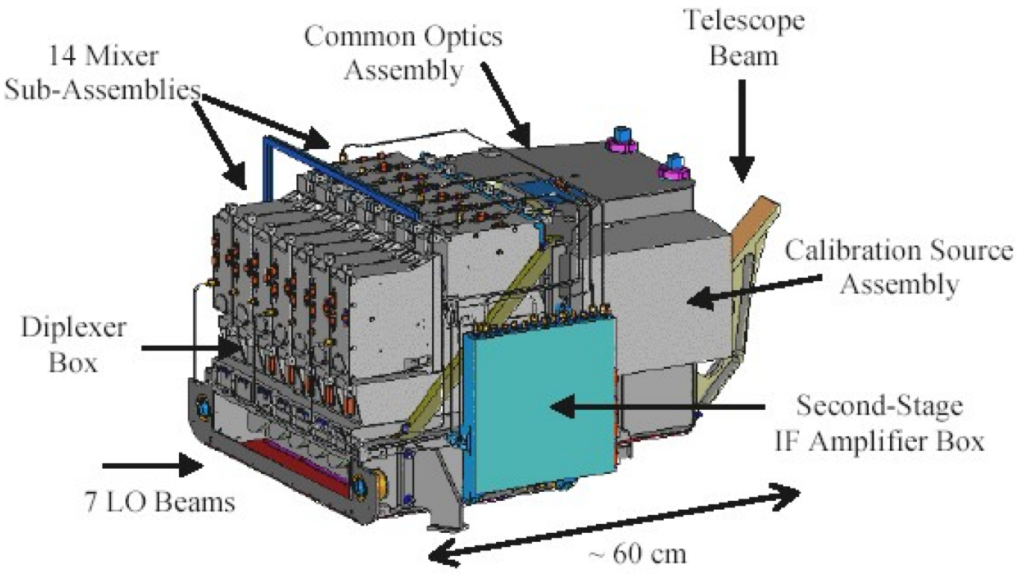

Figure 2.4: HIFI Focal Plane Unit (FPU).

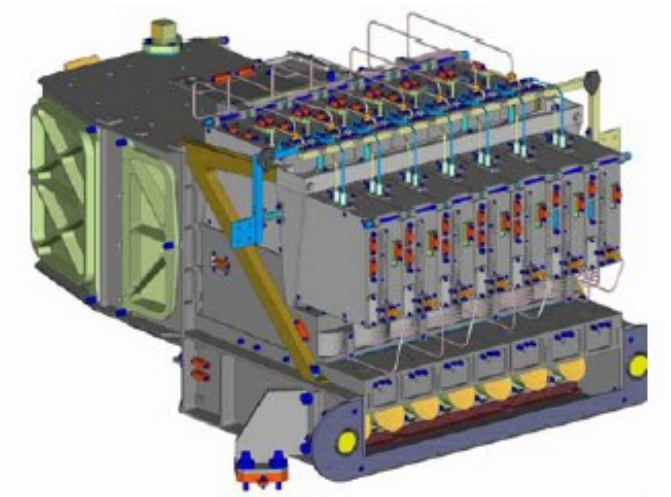

Figure 2.5: Back side of the HIFI FPU.

2.3 HIFI Focal Plane Unit

The HIFI Focal Plane Sub-System consisted of three hardware units: the Focal Plane Unit (FPU, see Jackson et al. 2002), which was located on the optical bench in the Herschel cryostat and depicted in Figure 2.4 and Figure 2.5; the Up-converter and 3-dB Coupler (described in Section2.5) which were contained in the satellite's service module – see the Herschel Mission & Satellite Overview Handbook for details on the service module; and the Focal Plane Control Unit (FCU), also contained in the satellite's service module). Additionally, the critical signal chain elements that together defined the instrument's sensitivity (the mixers, isolators in bands 1 to 5, and amplifiers, plus the IF up-converter used in Bands 6 and 7) together formed the HIFI Signal Chain.

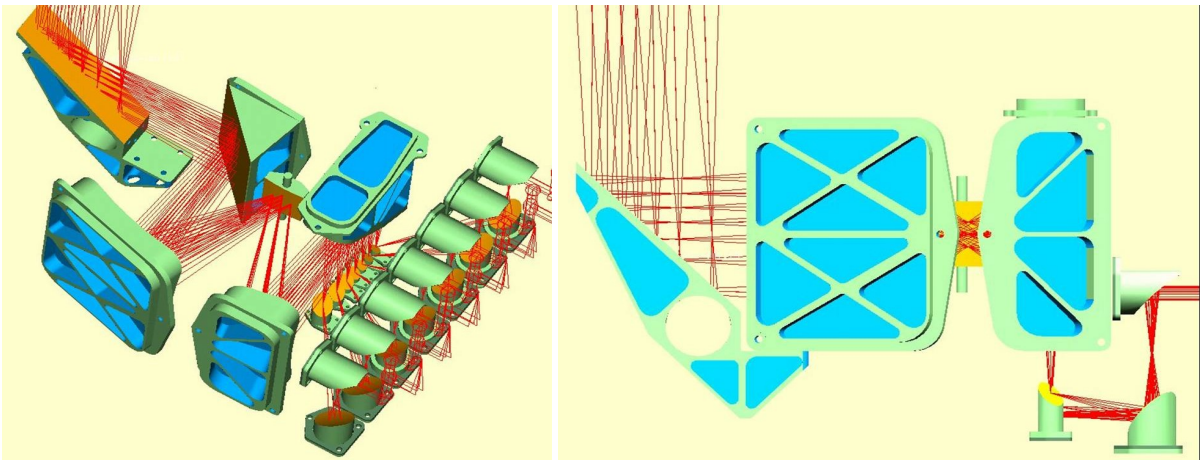

Figure 2.6: *Left*: HIFI telescope relay optics. *Right*: The channel splitting optics – as seen from the side with respect to the figure on the left.

2.3.1 The Common Optics Assembly

The Common Optics Assembly (COA) formed the basis of the FPU structure, and mounts directly on the Herschel optical bench.

The optics assembly relaid the instrument's 7 signal beams from the telescope's focal plane into a diplexer box. This was done with 6 common mirrors (the telescope relay optics), and 7 sets of 3 mirrors (the channel-splitting optics, see Figure 2.6). Together, these optics had three primary functions:

- They produced an image of the telescope secondary on the fourth mirror in the chain after the secondary (M6), enabling the implementation of a Focal Plane Chopper.
- They produced an image of the telescope focal plane on the first mirrors in the Channel-Splitting Optics, allowing the beams to be split by seven mirrors with different orientations.
- In each channel, they created an image of the telescope secondary within the beam combiner assembly (see Section 2.3.2). This image had a large Gaussian beam waist, to minimize diffraction losses, and a frequency independent size, to simplify visible-light alignment.

The seven local oscillator beams from the Local Oscillator Unit entered the FPU through windows in the cryostat. Using 7 sets of five mirrors, the Cold Local Oscillator Optics re-imaged the LO beam waists at the FPU input to waists in the diplexer box that matched those produced by the channel-splitting optics.

2.3.2 The Beam Combiner Assembly (Diplexer unit)

Within the beam combiner assembly, each of the 7 signal beams was combined with its corresponding local oscillator beam, creating two linearly polarized beams per channel (referred to as Horizontal, H, and Vertical, V, beams). Each of these 14 beams was then directed into a Mixer Sub-Assembly (MSA)

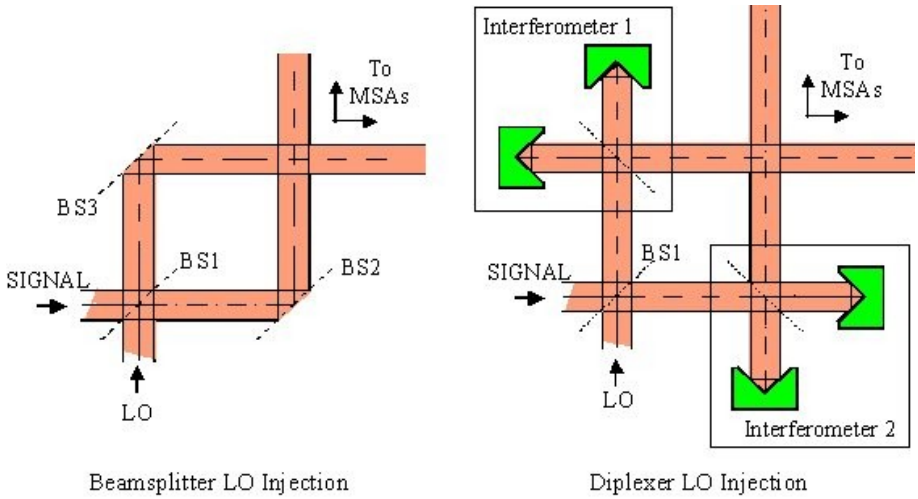

Figure 2.7: Beam-splitter and diplexer mixing with sample diplexer unit.

At low frequencies, where significant LO powers were available, the combining was done with polarizing beam-splitters. As seen in Figure 2.7, one beam-splitter was placed at the intersection of the LO and signal beams, creating two mixed beams (one contained the horizontally polarized signal beam and the vertically polarized LO beam, while the second contained the inverse). Each of the mixed beams then hit a second beam-splitter, which was oriented to reflect 90% of the signal power and 10% of the LO power (the remaining power was absorbed in a beam-dump).

At high frequencies, where LO power was scarcer, a Martin-Puplett diplexer was used for LO injection (see Figure 2.7). As in the beam-splitter channels, the first beam-splitter created two beams containing LO and signal power in orthogonal polarizations. However, in this case, the second beam-splitter was replaced with a polarizing Michelson interferometer that rotated the LO beam polarization relative to that of the signal beam, creating a linearly polarized output. In this manner, the coupling of both the LO and signal powers to the mixers was high (95%, or better), although diplexer scanning mechanisms were needed for frequency tuning.

2.3.3 The HIFI mixers

The mixers at the heart of the Focal Plane Unit largely determined the instrument's sensitivity. For this reason, the mixer technologies used in each band have been selected to yield the best possible sensitivity. In particular, a range of SIS mixer technologies are being used in the lowest 5 frequency bands (covering 480–1272 GHz; see Jackson et al. 2006, Jackson et al. 2005, Delorme et al. 2005 and Salez et al. 2003), while the top two bands (covering 1410–1910 GHz; see Cherednichenko et al. 2002, and Cherednichenko et al. 2005) incorporate HEB mixers.

Each of the 14 linearly polarized outputs from the diplexer/beam combiner box entered a Mixer Sub-Assembly (MSA – see Figure 2.8) that included:

- a set of three mirrors that focussed the optical beam into the mixer;
- a mixer unit where the incoming signal and LO signal were combined;

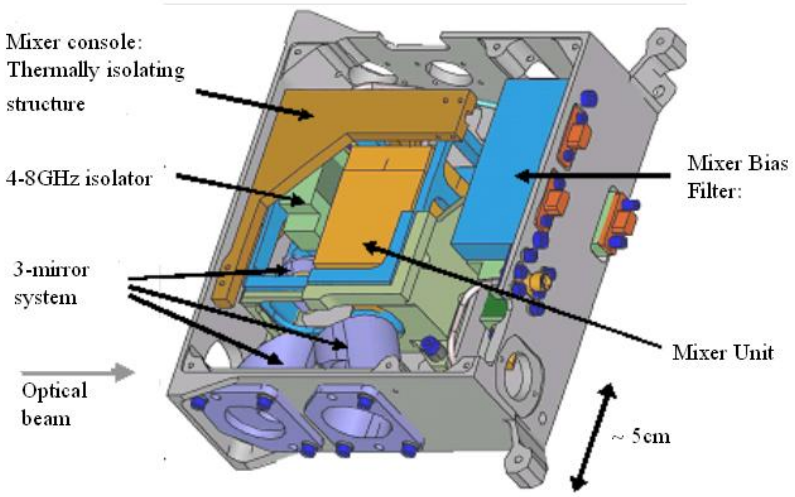

Figure 2.8: A HIFI mixer sub-assembly.

- a low-noise IF amplifier (plus two IF isolators - for bands 1 to 5 - that suppressed reflections in the cable between the mixer and the amplifier);
- low-frequency filtering for the mixers DC bias lines; and
- a mechanical structure that thermally isolated the mixer unit (at 2 K) from the FPU structure (at 10 K).

2.3.4 The Focal Plane Chopper

The Focal Plane Chopper (FPC) was the sixth mirror of the telescope relay optics (M6, see Figure 2.8). The chopper mirror was able to rotate (in one direction) around the centre of its optical surface. Tilting the chopper was equivalent to tilting the telescope secondary, which moved the beam on the sky. The primary use of the chopper was to steer the beam on the sky, and to redirect the instrument's optical beam into the on-board calibration sources.

The beam switch on the sky was a fixed parameter for the user, the beam switch being approximatively 3 arcmin on the sky. There were two chopper speed regimes available to the user, a "fast" chop (up to 4 Hz depending on the goal resolution, being faster for larger resolutions) and a "slow" chop (typically 0.125 Hz if all 4 spectrometers were used simultaneously, but twice as fast if two backends were switched off – e.g., when using the WBS only). The FPC was designed to have a settling time under 20 ms.

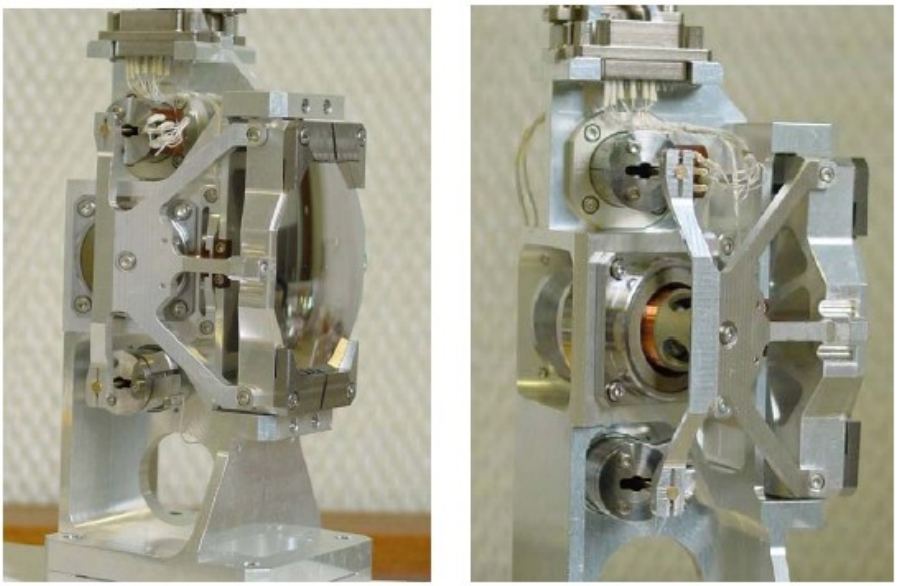

Figure 2.9: . The HIFI Focal Plane Chopper (FPC).

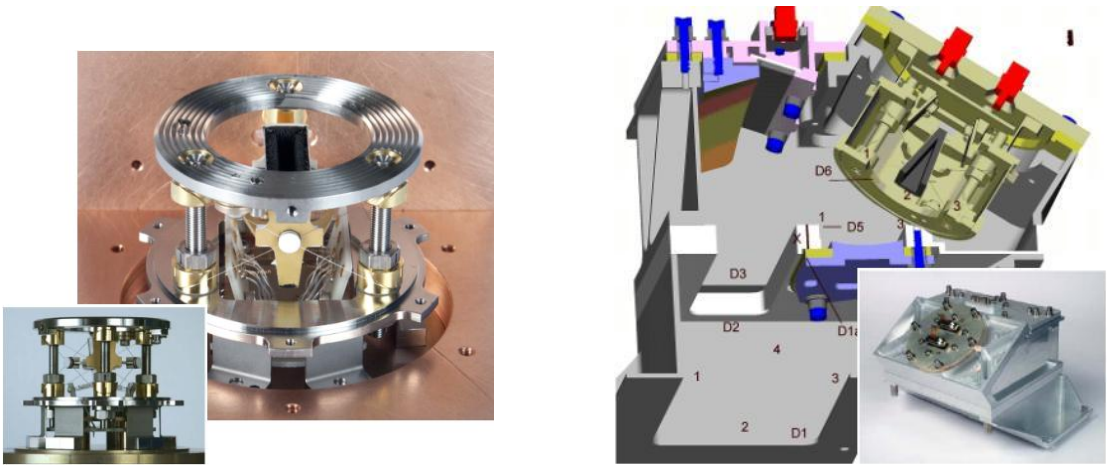

Figure 2.10: The HIFI Calibration Source Assembly (CSA).

2.3.5 The Calibration Source Assembly

Mounted on the side of the Common Optics Assembly, the Calibration Source Assembly (CSA, see Figure 2.10) included two blackbody signal loads that were used to calibrate the instrument's sensitivity (the first was a coated absorber at the FPU temperature around 10 K, while the second was a lightweight blackbody cavity that could be heated to 100 K), plus mirrors that focussed the FPU's optical beam into the loads. Temperature sensors were available to read out the actual temperature of both calibration loads. The HIFI optical beam was steered towards the calibration sources by the use of extreme positions of the Focal Plane Chopper.

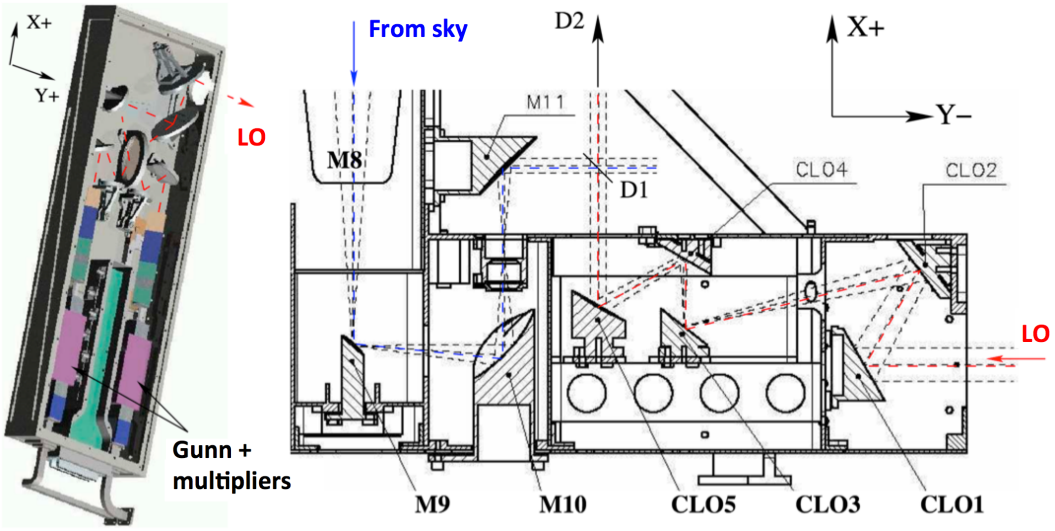

Figure 2.11: LOU assembly (left) and layout of the coupling optics between the LOU output and the FPU (right).

2.4 The HIFI Local Oscillator Unit

The Local Oscillator Unit (LOU) provided the LO signal to be combined with the sky signal at the mixer feed. The LOU contained the amplifier and multiplier assembly and the mechanical and optical structure necessary to feed the LO signal into the FPU. The Local Oscillator Source Unit (LSU) provided the first stage frequencies (total range 23.57–35.33 GHz) that was further up-converted by a triplet to the 71–106 GHz range. As such the HIFI LOU system was designed to directly use one single common synthesizer to drive the power amplifiers in all bands. Planar Schottky diodes were used for all the stages of the multiplier chains (Ward et al. 2003). These not only provided high power-handling capability and a wide bandwidth, but improved also considerably the re-producibility and stability, needed for a satellite project. In a next step, this signal was injected into the multiplier chain of the LOU power amplifiers in order to achieve the required RF converage of HIFI. These tuning ranges were achieved with broadband, high-power millimetre-wave amplifiers as input source for the varactor frequency multiplier chains. The demonstrated output powers of the amplifiers were sometimes in excess of 400 mW in the 75–100 GHz frequency range. Each mixer frequency band was covered by two chains in the corresponding LO assembly (LOA). The LO multiplication scheme is illustrated in Tab. 2.2 (from Pearson et al. 2000).

The LOU assembly was located outside of the *Herschel* cryostat and heated to its operational temperature range of 115–140 K by internal heaters. The LSU was located inside the SVM and its signal was fed to the LOU through waveguides also sitting in outer space. The LOAs contains several focussing mirrors to perform the optical beam transformation between the multiplier feed horn output and the FPU optical input (Fig. 2.11).

LO band	First stage (GHz)	Up-converter	Amplifier stage (GHz)	multiplier	LO frequency output (GHz)
1a	27.11–30.67	×3	80–92	×2×3	488–552
1b	31.44–35.17	×3	92–106	×2×3	566–633
2a	26.96–29.58	×3	80–92	×2×2×2	647-710
2b	31.17–33.04	×3	88–99	×2×2×2	724–793
3a	33.63–35.33	×3	92–106	×2×2×2	807–848
3b	23.94–26.47	×3	71–80	×2×2×3	852–953
4a	26.86–28.94	×3	80–92	×2×2×3	967–1042
4b	29.33–30.92	×3	88–99	×2×2×3	1056–1113
5a	30.67–35.33	×3	92–106	×2×2×3	1127–1178
5b	30.67–35.33	×3	92–106	×2×2×3	1192–1272
6a	29.70–33.20	×3	88–99	×2×2×2×2	1408–1575
6b	32.05–33.55	×3	92–106	×2×2×2×2	1528–1696
7a/7b	23.57–26.33	×3	71–80	×2×2×3×2	1710–1908

Table 2.2: Multiplier stages for bands 1 to 7 in the HIFI LO.

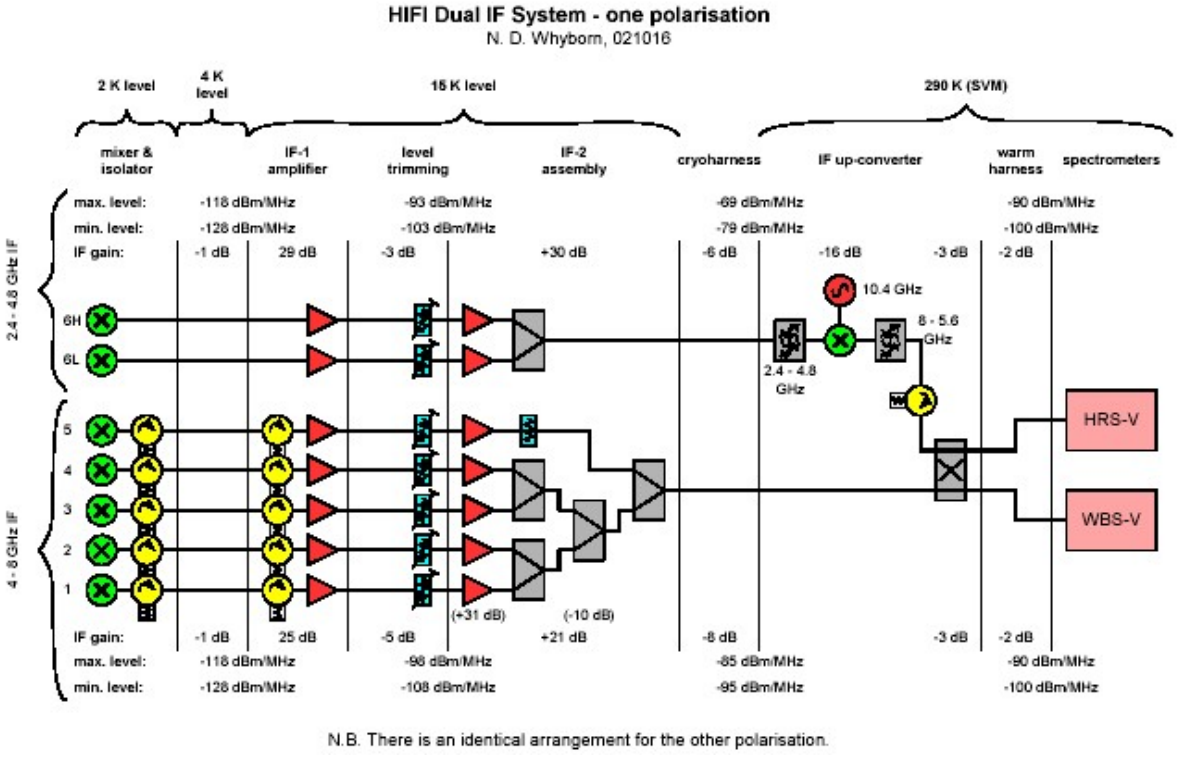

Figure 2.12: The HIFI signal chain.

2.5 The HIFI Signal Chain

As seen in Figure 2.12, the HIFI Signal Chain included the 14 Mixer Units, the First- and Second-Stage Amplifiers (in the FPU), the Upconverter and 3 dB Coupler (which were located in the satellite's service module), and the Isolators that were used to suppress reflections between the Mixer Units and the First-Stage Amplifiers.

In Bands 1 to 5 each Signal Chain consisted of a mixer, followed by two isolators and two amplifiers. For each polarisation, the outputs of these 5 independent chains were combined, so that only two cables were needed to carry the IF outputs from these 10 channels of the FPU to the service module.

The situation in Bands 6 and 7 was similar, except that isolators were not used (as they were not available for the 2.4–4.8 GHz IF band that was needed for the HEB mixers). Thus, the second pair of IF output cables from the FPU included the combined outputs of the two polarisations of Band 6 and 7. The other difference in the Band 6/7 Signal Chain was that an IF Up-converter was needed to transform the 2.4–4.8 GHz output of the FPU to 8–5.6 GHz, for compatibility with the spectrometers.

Within the "IF Up-converter" (in the service module), a 3 dB Coupler was also used to combine the Bands 1–5 and 6–7 outputs, so that each "polarisation" of the Wide-Band and High-Resolution Spectrometers was connected to all 7 bands by a single input cable (although a signal was only received from the active band).

2.6 The HIFI Spectrometers

The HIFI instrument provided an IF bandwidth of 4 GHz in all bands except for band 6 and band 7 (1430–1910 GHz) where only 2.4 GHz bandwidth was available. To sample this bandwidth, HIFI had 4 spectrometers. A Wide Band Spectrometer and High Resolution Spectrometer were available for each of the polarisations. All spectrometers could be used in parallel, although at fast data rates it was necessary to reduce how much was readout and stored since, at the highest data rates, the spectrometers provided data at a rate that was higher than the bandwidth available to HIFI on board the spacecraft.

The WBS was an Acousto-Optical Spectrometer (AOS) able to cover the full IF range available (4 GHz) at a single resolution (1.1 MHz). The HRS was an Auto-Correlator Spectrometer (ACS) with several possible resolutions from 0.125 to 1.00 MHz but with a variable bandwidth that could cover only portions of the available IF range. The HRS could be split up to allow the sampling of more than one part of the available IF range.

In the following two subsections, we describe the main workings of the two spectrometer types that were available to HIFI.

2.6.1 The Wide Band Spectrometer (WBS)

The WBS was based on two (vertical and horizontal polarisation) four channel Acousto-Optical Spectrometers (AOS; see Schieder et al. 2003) and included IF processing and data acquisition. To cover the 2×4–8 GHz (2×2.4–4.8 GHz for bands 6 and 7) input signals from the FPU, two complete spectrometers (horizontal + vertical polarisation) were used. For redundancy reasons both spectrometers were fully independent.

Each spectrometer received a pre-amplified and filtered IF-signal (4–8 GHz). After further amplification in the WBS electronics, the signal was split into four channels which provided the input frequency bands for the WBS optics (4×1.55–2.65 GHz; IF1 to IF4). The signal was further amplified and equalised (using variable attenuators), to compensate for non-uniform gain of the system, before being sent to two Bragg cells in the optics module of the WBS.

The other necessary input was to provide a frequency reference signal for the frequency calibration of WBS spectra. This was done using a 10 MHz reference signal from the Local Oscillator Source Unit (LSU), which was fed into the WBS to provide a "COMB" signal. The COMB signal in the WBS, with regular stable 100 MHz line spacing, could be connected for frequency calibration purposes or it could be disabled to provide a zero level measurement of the AOSs. The zero allowed more precise system temperature measurements to be made.

In the optics section of the WBS, the pre-processed IF-signal from the mixers was analysed using the acousto-optic technique. The IF-signal was fed into a Bragg cell via a transducer. The IF-signal then generated an acoustic wave pattern in the Bragg cell crystal. A laser beam which entered the Bragg cell was diffracted according to the acoustic wave pattern in the Bragg cell crystal. The diffracted laser light was afterwards detected by four linear CCDs with 2048 pixels each and each covering approximately 1GHz bandwidth. Four vertically aligned Bragg cells and CCD chains were necessary to cover the full 4 GHz IF bandwidth of HIFI.

The WBS electronic section had 4 analogue line receivers for the 4 CCD video signals. These signals were fed to 14 Bit analogue to digital converter with a conversion speed corresponding to less than 3 ms. The relatively high number of ADC-Bit was meant to keep differential non-linearity effects to a very low level. Overall non-linearity in the WBS was very low, less than 1%.

Continuous data taking was possible without dead time during data transfer, as long as the integration time was above one second – which was true for all standard operating modes of HIFI. Every 10 ms the collected photoelectrons in the CCD photodiodes were shifted into a register and clocked out serially. After integration completion, the data could be transferred while a new integration was started. Data was transmitted to the Instrument Control Unit (ICU) with 16 or 24 Bits through a serial interface with 250 kHz clock rate which was synchronous with the CCD read-out clock. Housekeeping data was provided through the same interface. A second serial interface was used for the command interface.

2.6.2 The High Resolution Spectrometer (HRS)

With the HRS, high resolution spectra were available from any part of the input IF bandwidth (4 GHz, or 2.4 GHz in band 6 or 7). The HRS was an Auto-Correlator Spectrometer (ACS) that could process simultaneously the 2 signals coming from each polarisation of the FPU. It was composed of two identical units: HRS-H and HRS-V, each of which included an IF processor, a Digital Autocorrelator Spectrometer and associated digital electronics, plus a DC/DC converter (not discussed here). The HRS provided capability to analyse 4 subbands per polarisation, placed anywhere in the 2.4 or 4 GHz wide input bands coming from the FPU. The two units of the HRS could be used to process the same 4 subband frequency ranges in each of the two polarisations provided by the FPU, thereby reducing the integration time and providing redundancy. Both units of the HRS operated at the same time and it was possible to look at either IF with each of the HRS spectrometers.

Mode	Number of band per polarisation × bandwidth (MHz)	Number of lags	Spectral Resolution (Hanning apodization) / Channel Spacing (kHz)
High Resolution	1 × 230	1 × 4080	125/64
Nominal Resolution	2 × 230	2 × 2040	250/125
Low Resolution	4 × 230	4 × 1020	500/250
Wide Band Resolution	8 × 230[a]	8 × 510	1000/500

Table 2.3: HRS resolution modes characteristics. [a] in the Wide Band mode, the bands had to go in non-separable pairs of 230 kHz, so in effect, there were 4 configurable bandwidths of 460 kHz each

Overview

In each HRS unit the ACS processed the signals coming from its associated IF (see Belgacem et al. 2004). Each 230 MHz band width input was digitised by a 2 bit / 3 level analogue to digital converter clocked at 490 MHz. The digital signals were analysed with a total of 4080 autocorrelation channels. It was possible to configure the HRS to provide 4 standard modes of operation as given in Table 2.3: wide band mode, low resolution mode, normal resolution mode, and high resolution mode. For example, in its nominal resolution mode, the HRS provided two subband spectra, each of which had a bandwidth of 230 MHz, each of which was covered by 2040 channels and had a spectral resolution of 250 kHz.

It was possible to set each subband frequency independently anywhere in the 4 GHz IF band range.

Two buffers were used, with selection synchronised with the chopper position by the ICU. The HRS had a maximum chopping frequency of 5 Hz. The data could be accumulated in each buffer up to a maximum of 1.95 seconds. The data readout duration was about 42 ms. Data could be read out from one buffer while data accumulation occured on the other.

Chapter 3

Observing with HIFI

3.1 Introduction

For HIFI, three Astronomical Observing Templates (AOTs) were offered to the community for observation planning via HSpot:

- **Single Point**, for observing science targets at one position on the sky;
- **Mapping**, for covering extended regions;
- **Spectral Scanning**, for surveying a single position on the sky over a continuous range of frequencies within one of HIFI's 14 LO subbands.

Each AOT could be used with different modes of operation, or Observing Modes, developed from principles found at many single-dish submillimeter telescopes on the ground for providing the widest practical range of reference measurements suited to the calibration accuracy and Signal-to-Noise Ratio (S/N) goals in a variety of astronomical settings. The development and optimisation of the Observing Modes were driven by HIFI's key science themes, emphasising very high velocity resolution of spectral lines occurring at frequencies inaccessible from the ground (de Graauw et al. 2008).

Users of HIFI data from the HSA should be familiar with how the observations were carried out in order to understand the quality and types of data products created in the HIFI pipeline. In the following section the basic design, usage, and limitations of the Modes are summarised. Calibration of the data in the HIFI pipeline is also briefly summarised for each mode, but for a full description including the mathematical formalisms of the calibrations, the reader should refer to Section 4.1. A description of the validations and performances of the Observing Modes is given in Section 5.5.

3.2 General Features of All HIFI Observing Modes

There are some general aspects that apply to all Modes:

1. Only one LO subband could be operated at any one time; i.e., the data products from the SPG for every observation are associated with a single LO frequency (for Point or Mapping AOTs)

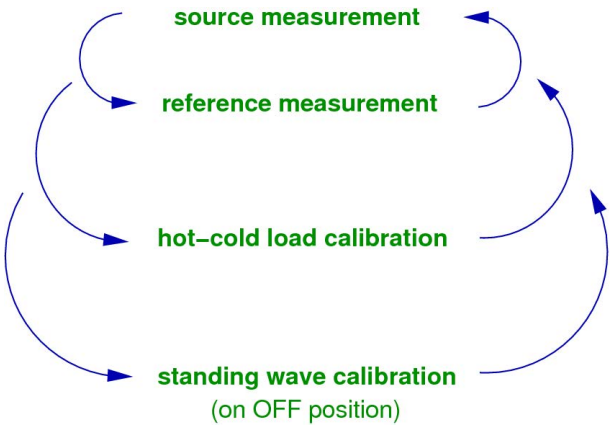

Figure 3.1: Schematic approach of the timing loops in HIFI Observing Modes.

or range of frequencies (for Spectral Scans) within the same LO subband. In HIPE it is possible to interactively combine or merge observations taken at different LO frequencies without restriction to the LO subband. Some UPDPs and HPDPs available in the HSA (Section 6.2) have been created this way.

2. The instrument and telescope setup of each observation is contained in an Astronomical Observing Request (AOR) created by the original Observer in HSpot. The AORs created in the Key Programmes and Open Time Programmes can be retrieved and visualised in HSpot[1]. Visualisation is often helpful to check the location on the sky where reference measurements were taken (e.g., for possible contamination), beam size references and focal plane layout, the location of the targeted spectral line(s) in the upper or lower sideband using the HIFI Frequency Editor, and so on.

3. Integration times have been optimised by the HIFI time estimator in HSpot according to the original user's observing time goal or noise level goal in the AOR. The resulting noise performances, both predicted and measured in the HIFI pipeline, are found as metadata in the Observation Context. Further explanations are given immediately below and in Section 5.5.

In the descriptions below, the roles of the Observer and user are often referred to, where the Observer has been the creator of the AOR that was carried out, and the user can be both the Observer and the astronomer with an interest in HIFI data products in the HSA.

The Role of Spectral Resolution and Noise Goals

A major function of all HIFI Observing Modes is the inclusion of loops between observations of the astronomical source and measurements of reference and calibration sources in order to account for instrumental drifts. The calibration sources could be the internal hot (≈100 K) and cold (≈10 K) thermal loads that provide intensity calibration, and a comb source that provides the frequency calibration of the WBS. Reference measurements were provided from chopped observations of the sky for DBS modes, a thermal load in Load Chop modes, or an offset frequency in the Frequency Switch modes.

[1] See the HSC's webpage for HSpot download and installation instructions, http://www.cosmos.esa.int/web/herschel/ao2-tools#HSpot

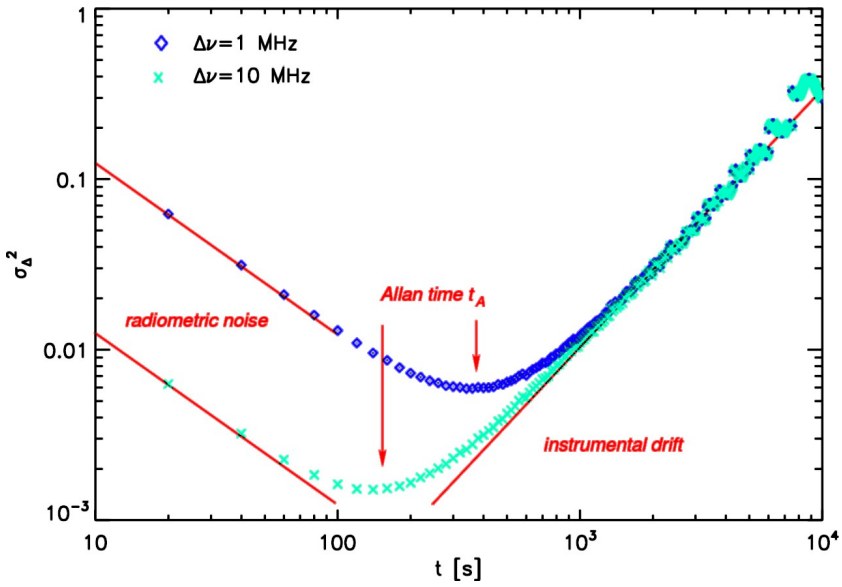

Figure 3.2: Variation of measured radiometric and drift noise with observing time at a representative frequency over fluctuation bandwiths of 1 and 10 MHz. The Allan time is defined by the crossing ponit between radiometric and drift noise, and depends on the spectral smoothing width.

A reference measurement of a line-free position on the sky, designated by the Observer, also provided standing wave corrections for optical path variations when using the internal chopper mirror for sky and internal reference and calibration measurements. A schematic illustration of the looping is shown in Figure 3.1.

The timing of these loops depended strongly on (1) stabilty (Allan) times and system noise temperatures $T_{\rm sys}$, and (2) the noise goals set by the Observer in HSpot at the frequency or frequencies of interest[2] These noise goals were driven in large part by the spectral resolutions that were appropriate to the science goals of the observations. In other words, obtaining the observations with baseline noise goals at the "native" resolutions of the WBS (1.1 MHz) and/or HRS (0.125 to 1.00 MHz) could be very expensive in observing time, and thus the Observer could set noise goals at lower resolutions tuned to expectations of spectral line strengths and widths, and S/N requirements. The resolutions were set by the Observer in HSpot by minimum and maximum fluctuation bandwidths, appearing in data products from the HSA as `noiseMinWidth` and `noiseMaxWidth` in the Observation Context. The value of `noiseMinWdith` drove the timing of the calibration loops, since Allan variances at any frequency is not a constant with spectral resolution; see Figure 3.2. The detailed concepts of a "unified" Allan variance approach in Observing Mode design are given by Ossenkopf (2008).

The optimisation of the timing loops, under the previous mentioned boundary conditions of system response and the observing setup and noise goals set by the Observer, was achieved during time estimation in HSpot with a sequencing engine. The Sequencer's main task was to minimize the observing cost function specific to the observation, using inputs from the AOT logic for the specific Observing Mode and setup, and outputting a set so-called sequence parameters such as data readout interval, thermal load period, OFF reference cycling, and so on. These parameters found at the lowest

[2]The timing is optimzed at the requested LO frequency in Point and Mapping AOT observations, and at a particular reference frequency determined in the AOT logic for Spectral Scan observations.

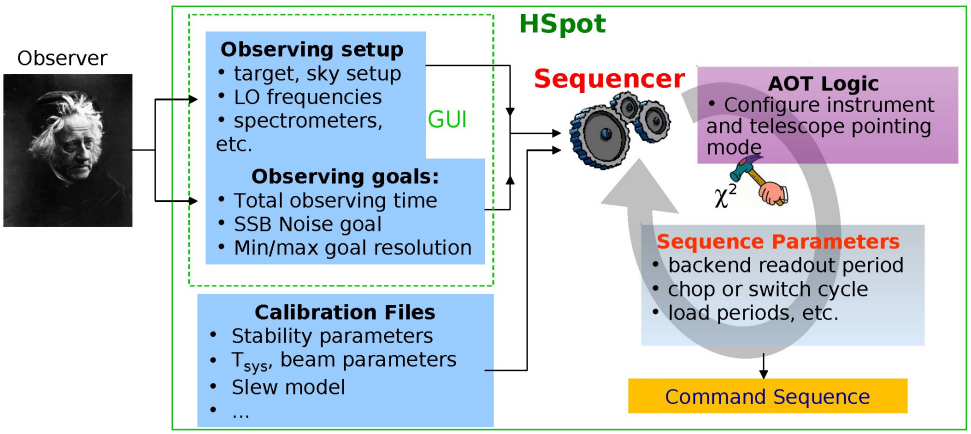

Figure 3.3: Cartoon illustration of the flow from Observer inputs via the HSpot front-end GUI (dashed green box) and input uplink table parameters (blue boxes) to the AOT logic (magenta box) which builds the command sequence based on parameters which have been optimised by a sequencing engine to minimise the observing cost function χ^2.

value of χ^2 = (total noise)^2 $\times$ $t_{\rm obs}$+ penalties, where (total noise)^2 = (radiometric noise)^2 + (drift noise)^2, were used as inputs to the time esimation code of the AOT logic to project noise performances. The Observer could iterate on these estimations until settling on the final version of the AOR. A cartoon of this process is shown in Figure 3.3.

It is important for users of the HSA to understand these aspects about the noise performances in the HIFI data products:

1. The Observing Mode references and calibrations were taken in loops whose timing depended on system performances at a given LO frequency and the Observer's noise goals at the minimum fluctuation bandwidth set in HSpot;

2. The spectral data products in the HSA have been processed at the native resolution of the spectrometers, but noise levels appropriate to the Observer's science goals for that observation are acheived only after smoothing to the spectral resolutions set by the Observer.

3. Spectral products from the standard HIFI pipeline are not smoothed, but there are steps carried out in the pipeline to process the Level 2 spectra to a form that can be compared to the HSpot noise predictions (i.e., spectral smoothing at the specified bandwidths, combining H and V polarizations, and so on). The pipeline records the measured noise levels at both native and smoothed resolutions, providing these as metadata in the Observation Context. Performances and comparisons are described further in Section 5.5.

3.3 Description of the HIFI Observing Modes

The Observing Modes offered with the three HIFI AOTs differ from each other mainly in the selection of the reference measurements during the course of observing. All observations consisted of ON-source measurements, reference measurements and a set of calibration measurements that have been

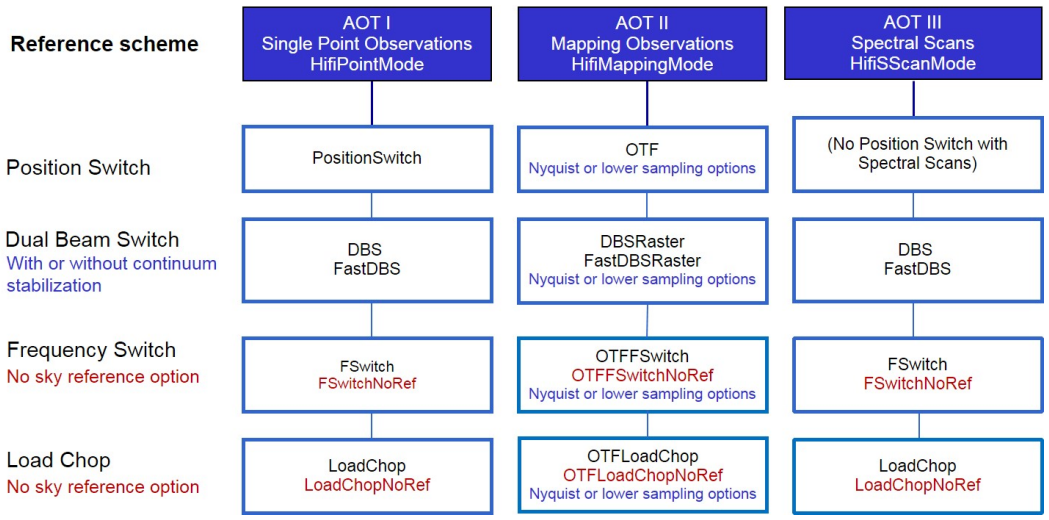

Figure 3.4: Overview of available AOT Observing Modes.

used in the HIFI pipeline to fully calibrate the spectra in frequency and intensity. Observing Mode design was intended to supply an optimum balance between observing efficiency and self-contained calibrations timed by instrumental performance and stability metrics.

The available Observing Modes and their relation to the AOTs are shown as a chart in Figure 3.4. When searching or browsing the HSA for HIFI observations, the Observing Mode name is returned after a query, and it can also be used as an input search item. They also appear in metadata with the data products, e.g. `instMode`. The names are prefixed by the AOT name as indicated in each AOT box, and appended with the Observing Mode designation as given in the Reference Scheme boxes. A Spectral Scan observation which used the internal thermal load chopping scheme has the mode name `HifiSScanModeLoadChop`, for example. This chart does not include corresponding calibration versions of the Observing Modes, the so-called Proc Modes which have been used in some calibration activities and often contain scientifically useful data available through the HSA (see Section 3.3.4).

Figure 3.4 indicates the main options that were available with the Reference Schemes. The Dual Beam Switch (DBS) modes employing an internal chopper mirror were available with two ranges of chop rate, a slow (or normal) DBS mode, and a fast DBS mode. Both versions had a continuum stabilisation option in HSpot. The schemes using internal reference measurements, the Load Chop and Frequency Switch modes, were normally employed with blank sky reference measurements for standing wave corrections, but had the option to skip the sky measurements, defining the "`NoRef`" modes. The usage of these modes and their options in various astronomical settings is summarised in the following sections. Each Observing Mode uses a somewhat different scheme for the data processing including the intensity and frequency calibration depending on how the reference measurements were obtained while observing. Thus the noise level and the drift contribution to the total data uncertainty of the calibrated data obtained from one of the AOTs depend critically on the Observing Mode.

The following subsections provide simplified descriptions of the scientific motives, typical usage, user options, advantages and disadvantages, and data calibration (in descriptive terms) of the AOTs and modes, as they were presented to the community for observation planning in HSpot[3]. Specific

[3]These historical references are the HIFI Observers' Manual, and the more technical but less complete HIFI Observing Modes Description Document, Ossenkopf & Morris (2005)

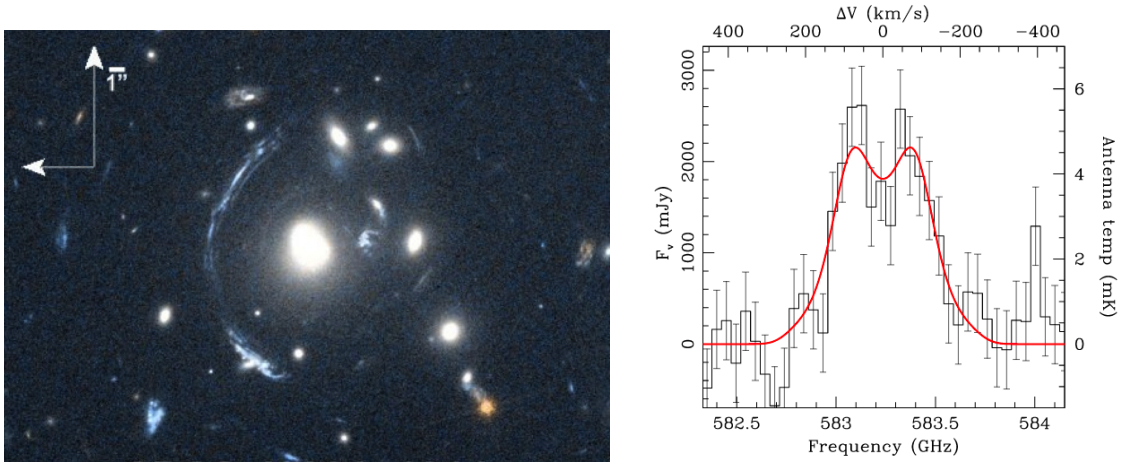

Figure 3.5: A HIFI Point AOT example. The optically lensed galaxy S0901 at a distance $z \approx 2$ observed with HST/WFC3 (left), and a HIFI observation of emitted C^+ redshfited from its rest frequency of 1900.537 GHz (right). The HIFI data were taken with the WBS smoothed to 25 km s^{-1} using the Point Fast DBS Observing Mode. [Credits: (left) NASA/STScI, the Master Lens Database, Moustakas et al. (2014); (right) Rhoads et al. (2014).]

recommendations were provided by the HIFI ICC to many KP and OT programme investigators for optimum mode selection and setup, but were not in any way restrictive. On the contrary, each Observing Mode was designed with maximum practical flexibility in mind, and many sources were observed with several modes, while the same Observing Mode could be used to observe very diverse celestial targets.

3.3.1 Modes of the Single Point AOT

For observing at a fixed position on the sky at a fixed LO setting that provides two sideband ranges of sky frequencies simultaneously, four modes were developed. The mode recommended to the Observer obviously depended on the science objectives and the properties of the target. A point source well clear of any diffuse cloud emission at the targeted line frequency is likely to have been observed with a DBS mode in which reference OFF-source positions are taken at a fixed throw of $3'$ and position angle (determined by the telescope's position angle at the time of the observation) on either side of target object. Sources embedded within extended molecular clouds, on the other hand, may have been observed with a load chop or frequency switch mode, using a sky reference measurement if blank sky was available within an allowed angular distance (up to 2°), or skipped if no blank sky was available or else for efficiency reasons.

Position Switch (`HifiPointModePositionSwitch`)

This was the simplest Observing Mode for HIFI, in which the single pixel beam of the telescope is pointed alternately at a target (ON) position then a reference blank sky (OFF) position, within 2° of the target position. Observing was done at a single LO frequency set by the Observer so that the targeted spectral line falls in one or the other sideband, at the spectral resolution of the chosen backend spectrometer(s). Data taken at the OFF position provide the underlying system background

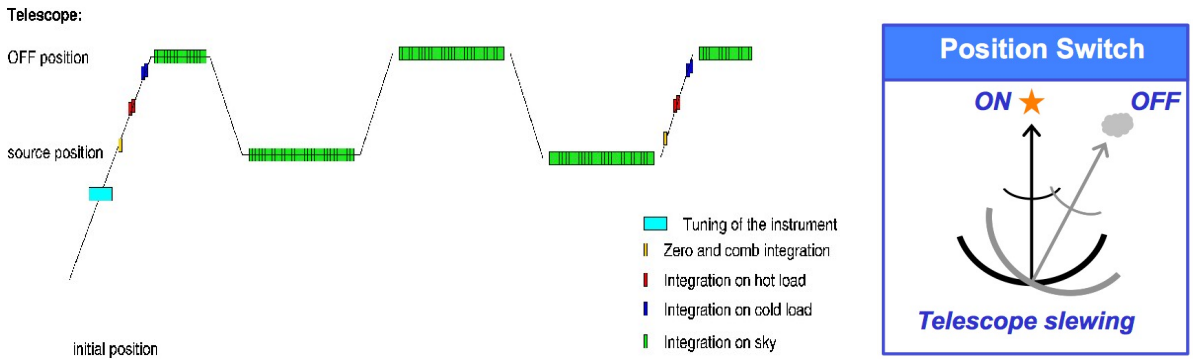

Figure 3.6: Schematic timeline of the HIFI Point Position Switch AOT.

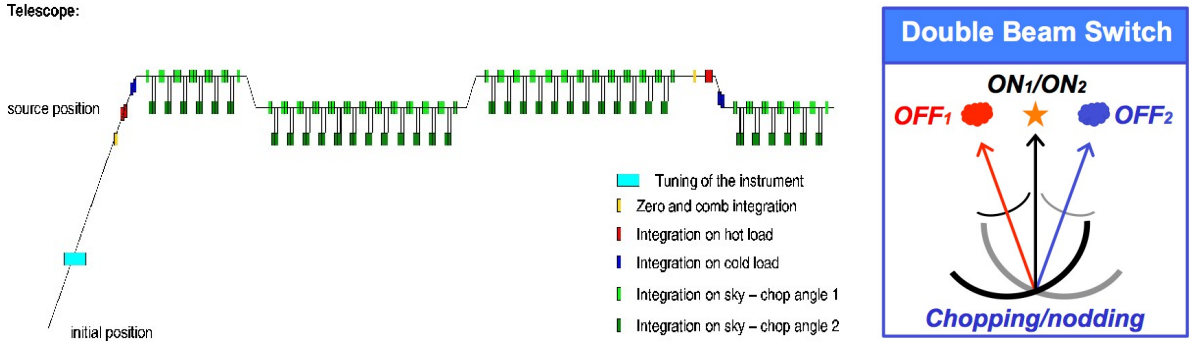

Figure 3.7: Timeline of the HIFI Point Dual Beam Switch AOT.

that is removed in pipeline processing by a simple subtraction. The reference position was sampled with sufficient frequency that detector drifts could be adequately compensated. The switching rate was calculated automatically in the AOT logic based on knowledge of the instrument stability time, and internal calibrations (e.g., frequency calibration) were taken during initial slews to the target and telescope movements between ON and OFF positions. This total power mode was most often unable to fully compensate for drifts at LO frequencies with short Allan times, particularly in Bands 6 and 7, leaving both optical and electrical standing wave residuals in the final spectra from the SPG. These are treatable with interactive tools in HIPE (Section 6.4.2).

Schematically, a timeline for this mode and the associated telescope slew phases is represented in Figure 3.6.

Dual Beam Switch (`HifiPointModeDBS` and `HifiPointModeFastDBS`)

In this mode a combination of an internal chopper mirror and telescope nodding provided the motion to reference OFF positions on the sky. The reference OFF position was at a fixed $3'$ throw on either side of the ON target position, at a position angle set by the observing date. Neither the throw distance nor position angle could be adjusted in the AOR by the Observer, although the position angle[4] could be specified by requesting a fixed observation date when the telescope roll angle could be predicted. In practice a "patchy" region within $3'$ of the target was rarely known well enough for the Observer to

[4]In DBS mode the throw angle is actually two angles, ϕ and $\phi + \pi$.

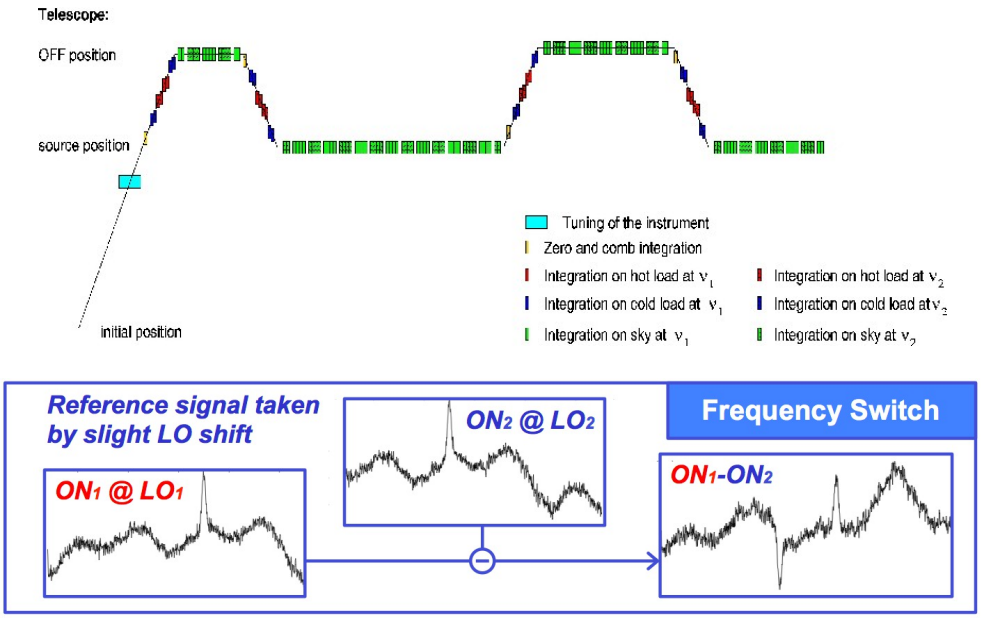

Figure 3.8: Timeline of the HIFI Point Frequency Switch AOT.

request a fixed time observation to avoid line contamination. Moving the internal mirror changes the light path for the incoming waves, subjecting the measurements to the possibility of residual standing waves, so nodding the telescope in such a way that the source appears in both chop positions was designed to completely eliminate the impact of standing wave differences in the two light paths.

The low dead time in moving a small distance with the telescope and in the internal chopper motion made this one of the most efficient modes offered, and not surprisingly the most popular. A diagram showing the telescope and chopper positions is shown in Figure 3.7, together with a schematic of the observation timeline.

Two chopper speeds were available to the user, a slow chop mode with typically 0.125 Hz chopper speed, and a fast chop with a frequency up to 4 Hz. The faster chop was intended to compensate for instrumental drifts at low frequency resolutions that would otherwise (at the slow chop rate) lead to baseline distortion and increased intensity uncertainties. For observations where accurate continuum level measurements were important, e.g., for some absorption-line studies, then a separate calibration scheme could be selected by the user in HSpot to provide a stable continuum. The continuum stability timing could significantly reduce the efficiency of the observation, particularly at the highest LO frequencies in Bands 6 and 7.

Frequency Switch (`HifiPointModeFSwitch` and `HifiPointModeFSwitchNoRef`)

In this mode integrations are performed at two LO frequencies, first at the target frequency requested by user followed by a slightly offset frequency that causes the spectral ranges in the upper and lower sidebands to be commensurately shifted. A schematic timeline of frequency switch measurements is shown in Figure 3.8. The two phases are differenced in the pipeline to correct for baseline artefacts. The frequency shift or “throw” has been selected from three options in each band by the user in HSpot (see Table 3.3.1 for the frequency throws that were offered), optimised to cancel out constructive

Band	Large negative (MHz)	Small negative (MHz)	Small positive (MHz)	Large positive (MHz)
1a to 4b	-288	-94	+94	+288
5a to 5b	-300	-98	+98	+300
6a to 7b	-192	-96	+96	+192

Table 3.1: Frequency throws offered depending of the HIFI band.

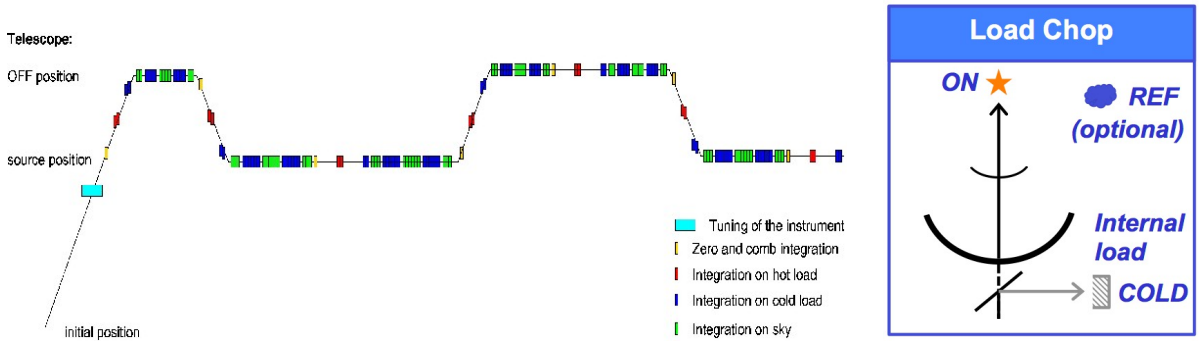

Figure 3.9: Timeline of the HIFI Point Load Chop AOT.

optical standing waves, and chosen by the user so that the spectral line(s) of interest remain observable at both LO frequencies without overlapping profiles in the two phases. This implies that the user expected an uncomplicated spectrum and that the line widths at the baseline (the Full Widths at Zero Intensity) would be less than half the distance to the selected throw, i.e., observations of very broad lines were not suited for this mode. Since the offset frequency also contains lines of interest in one of the sidebands (it is also taken ON-source), these appear as inverted lines in the difference spectrum.

The mode was very efficient, since target emission lines were observed at both frequency positions. However, the system response is likely to differ between the two frequency settings and simple spectrum arithmetic was not always anticipated to result in clean baseline removal, particularly at LO frequencies with short stability times. Residual artefacts could be mitigated by observing both frequencies also at a sky OFF position, which functioned the same as a Position Switch. The sky reference option could be skipped (although discouraged) in the `NoRef` version of the mode if no line-free region within 2° was available, or the Observer wished to trade this correction for observing efficiency. In this case the user has judged that baseline artefacts were not detrimental to the science, or that the standing waves and baseline drift could be acceptably corrected interactively.

It should be noted that Frequency-switched observations need to be ”folded” in order to fully benefit from the noise improvement implied by the observing mode. This folding is only done in the processing from Level 2 to Level 2.5 products (see Section 6.1).

Load Chop (`HifiPointModeLoadChop` **and** `HifiPointModeLoadChopNoRef`)

In this reference scheme, an internal thermal calibration source is used as a reference to correct short term changes in instrument behaviour. The HIFI chopping mirror alternates between the target on the sky and the internal cold (∼10 K) source with a typical period of a few seconds. The general timeline is illustrated by the schematic picture in Figure 3.9. This mode has been particularly useful when

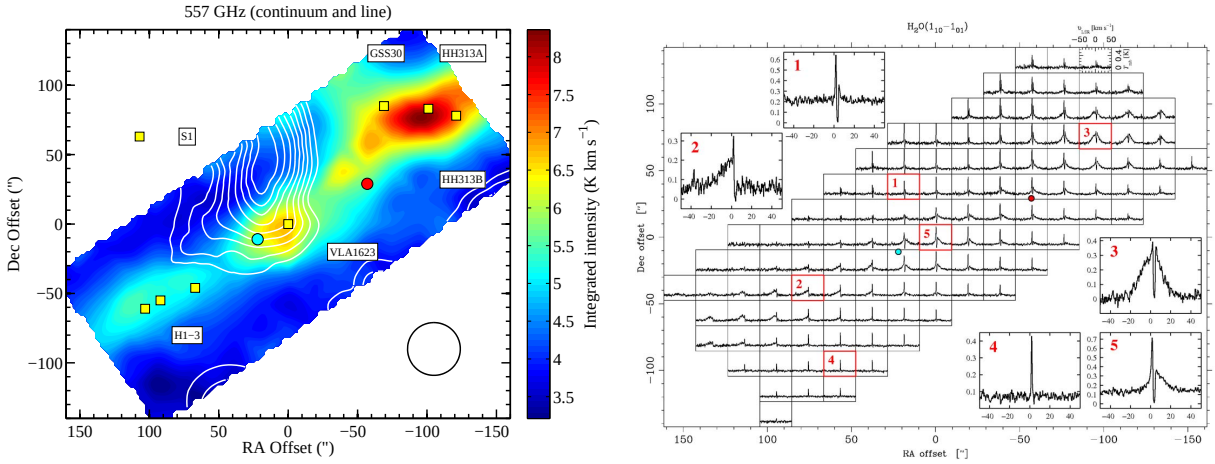

Figure 3.10: A HIFI Mapping AOT example. On the left, in colour an integrated intensity map of H_2O $1_{10} - 1_{01}$ emission centered near the Class 0 YSO outflow source VLA 1623 (blue dot), integrated over -40 to $+55$ km s^{-1}, with contours of the 557 GHz continuum overlaid. On the right, spectral extractions on the gridded map which each plot on a T_{mb} intensity scale. The data were taken with WBS using the OTF Position Switch Observing Mode. [Credit: Bjerkeli et al. (2012).]

there are no spectral line-free regions near the target that could otherwise be used as reference in either Position Switch or DBS modes, or where Frequency Switching cannot be used due to the spectral complexity of the source or in HEB bands where the latter did not offer very good performance. As such this has for instance been the workhorse mode for projects dedicated to the study of the [CII] line emission. Since the optical path differs between source and internal reference, a residual standing wave may remain in the `NoRef` mode. Additional measurements of an OFF reference position (blank at the targeted line frequency) reduces baseline ripple. The Load Chop scheme has proven robust but had relatively high dead times. The total time spent on the OFF position depended on the frequency resolution needed to describe the baseline ripple which could be considerably smaller than the integration time spent on the source in some observations.

3.3.2 Modes of the Mappping AOT

Spectral mapping with HIFI could be done either "On-the-Fly" (OTF), in which the instrument was continuously taking data with the telescope in a line scanning slew mode, or as a raster map on a user-defined grid using one of the DBS modes. OTF mapping has been offered with standard Position Switching, Load Chop (with or without Position Switch to a sky reference), and Frequency Switching (with or without sky reference) for the calibration schemes. Similarly, like the Point AOT, the DBS raster modes were offered with normal and fast chop speed ranges and the continuum stabilisation option. All maps could be taken at either Nyquist sampling, or at fixed spacings of the beam centres by 10″, 20″, 40″, or $0.5 \times \theta_b(\nu)$ (the so-called "Half Beam" option), where $\theta_b(\nu)$ is the Half Power Beam Width (HPBW) at frequency ν.

There are two historical notes to the mapping modes as they were first offered in HSpot[5]:

[5]Further information is found in Section 5.5

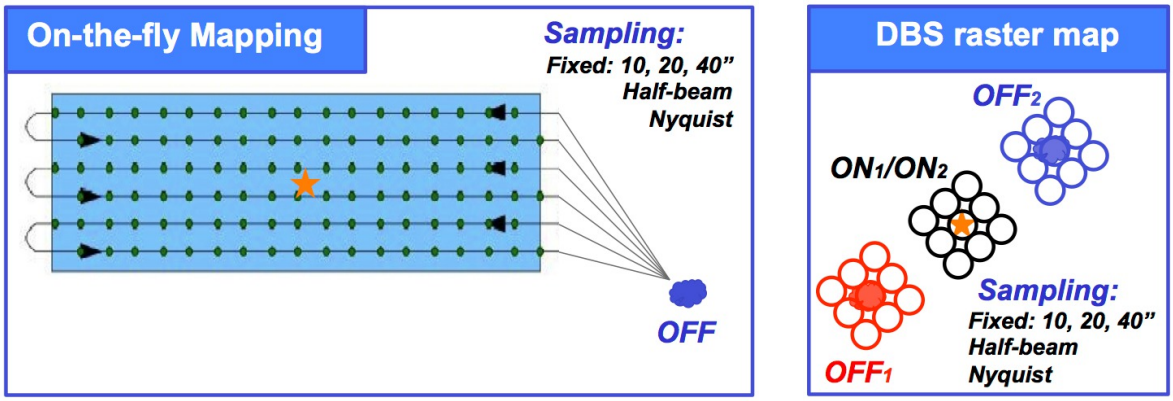

Figure 3.11: Left: Timeline of the path on the sky for a typical OTF map with Position Switch to an OFF position. The green dots represent regular data readouts of the spectrometers. The OFF position is returned to at the end of an integer number of scan lines. Right: Schematic representation of a 3×3 DBS raster map.

1. A variant of the DBS raster map, a 5-point raster cross map, was tested early in the mission (so users of the HSA may find a few `HifiMappingModeDBSCross` observations) and intended mainly for compact objects with poorly known positions or observations of solar system objects which had not yet been proven in terms of path planning and telescope tracking performance, but ultimately the DBS cross mode was never released to the community.

2. For the first part of the mission, until Operational Day (OD) 490, maps which requested Nyquist sampling were actually carried out with Half Beam spacing. The affected observations were all science and calibration programme DBS raster maps, and only calibration OTF maps (which had not yet been released for scheduling of KP AORs). In maps taken after OD 490, Nyquist spacing was correctly defined as $\theta_b(\nu)/2.4$, and Half Beam sampling was introduced in HSpot as an option.

OTF Maps with Position Switch Reference (`HifiMappingModeOTF`)

OTF mapping with Position Switch reference could be a very efficient observing mode since sky reference measurements at a user-defined OFF position could be shared with several ON-target scan lines (depending on timing constraints set by instrument stability, noise goals and overheads such as slew times), and the ON-source data were continuously read out while the telescope scanned over the target area to be mapped. The mode is relatively insensitive to optical standing wave problems because only one optical path is used. Electrical standing waves (ESW) could be more difficult to correct with the reference measurements self-contained in the observation due to the short Allan times in HEB Bands 6 and 7 (Section 5.2).

After each scan line, the telescope switched scanning direction so that the next line was scanned in the opposite direction. Normally integrations were stopped during these turns, but taking serendipity data was possible and did not interfere with the optimum start and stop times for the integration within each line. After a period determined by the system stability, the ON-target scanning was interrupted for the reference measurement at the OFF position. Reference measurements were only taken after

a complete number of scan legs. Changes in the instrumental sensitivity were measured in the frame of a fourth loop using the known difference in the thermal field between the internal hot and cold loads. This load calibration was performed during slews to the OFF position. The efficiency of the mode profited from high scan velocities, and OTF maps were typically observed in a series of coverages or cycles, adding up to the required total integration time per source position. The timeline of the observation consisted of telescope motions across the map and between the map and the OFF position, integrations by the instrument during the different phases and interleaving load calibration measurements. The left panel of Figure 3.11 shows the timeline for the OTF mapping mode.

In the example shown in Figure 3.11, the OFF position is visited three times within one coverage of the whole map, after every two scanned lines. A series of subsequent coverages could be performed in the general case, with an extra OFF measurement at the end of the observation guaranteeing a complete enclosing of the source observations by OFF observations. The green dots symbolise the points where the spectrometers are read out. The integration starts as soon as the telescope enters the blue area of the map. Note the change of the scanning direction after each row.

In the event that no nearby line-free position could be located within a typical angular distance of less than 1–2 degrees, a nearby "auxiliary" OFF could be defined and some programmes utilised a separate Point AOT observation at the same LO frequency (particularly for [C II] and [N II] observations) to characterise the emission at that position.

OTF Maps with Frequency Switching (`HifiMappingModeFSwitchOTF` and `HifiMappingModeFSwitchOTFNoRef`)

In OTF mapping with Frequency Switching a scan map is carried out in the same manner as OTF mapping with Position Switching regarding the line scanning and slewing to an OFF sky reference position after an integral number of scan lines. In the `NoRef` version the sky reference measurement is skipped, although this was generally discouraged. During line scanning, a small switch of the LO frequency is introduced as described for a Point Frequency Switch observation, and this reference spectrum may be used to calibrate neighbouring map points taken both before and after the reference measurement. The effect of introducing a reference measurement while scanning is to leave a gap in sky coverage at those positions. To compensate for this effect, the AOT logic optimises the spacing of readouts along each scan line to provide an average separation between readout points that matches the requested sampling density as closely as possible, balanced by requested map size.

Overall this observing mode was very scarcely used, and almost exclusively in the framework of calibration and AOT validation activities.

OTF Maps with Load Chop (`HifiMappingModeLoadChopOTF` and `HifiMappingModeLoadChopOTFNoRef`)

OTF mapping with Load Chop operated similar to OTF with Frequency Switching, which is to say that they both operated as a standard OTF map with Position Switch to a sky reference chosen by the Observer in HSpot. Like the Frequency Switch mode, a reference measurement could be taken periodically over each scan line. For Load Chop the reference source is the cold ($\sim$10 K) thermal load, which is measured as an OFF using the focal plane chopper mirror during OTF line scanning. Load chopping during OTF mapping was a good mode for situations where fast chopping to references has been needed for optimum stability, such as in Bands 6 and 7 or specific LO frequencies with low

power or short Allan times. As such it was the prime mapping mode for large scale mapping of strong lines such as high-J CO or [CII] (e.g. Goicoechea et al. 2015).

In bands 6 and 7, where most of the baseline artefacts arise from the Electrical Standing Wave (ESW, Section 5.3.1), it is an alternative option to actually disable the OFF position subtraction when line contamination is considered an issue, and treat the ESW on data only taken ON-target. See Section 6.4.1 for more details.

DBS Raster Maps (`HifiMappingModeDBSRaster` and `HifiMappingModeFastDBSRaster`)

In HIFI's raster maps a grid of fixed positions are measured using a Dual Beam Switch method. The mode works for each grid position in a similar way as for the Point DBS mode. However, the DBS may be applied at each map point individually, i.e., in a nodding-in-raster fashion that would mimic the Point DBS mode, or to the entire set of map points in a nodding-of-raster fashion. The right panel of Figure 3.11 shows a schematic layout of a 3×3 raster map. The choice was not made by the Observer but in the AOT logic according to timing constraints. To the archive user the difference of these pointing modes should not be apparent in the data, the products are indistinguishable. Like the Point DBS mode, the beam switching sampled the sky at 3′ on either side of each ON-target grid position, with a telescope movement occurring between changes to the beam switch either side of the grid position. With this setup, only small maps were recommended for this mode since maps larger than 3′ in the chopper throw direction could result in chopping of nearby emission or absorption features at the frequency of interest.

The raster map mode was the prime mapping mode used to characterise the HIFI beams. Indeed, such measurements being performed on strong continuum sources (Mars in particular), the continuum stability was of prime importance and a mapping mode combined with a fast referencing scheme was the best option for this.

3.3.3 Modes of the Spectral Scan AOT

Spectral scans consisted of a series of observations of a single target on the sky (either fixed or tracked as a moving target) at several frequencies using the WBS as main backend; see Figure 3.12 for a schematic representation. Depending on data rates, HRS data may also have been collected in a serendipity mode at no extra charge to observing time. After data processing, the result of such an observation is a continuous single-sideband (SSB) spectrum for the selected position covering the selected frequency range (Fig. 3.13). While observing, LO tuning was advanced in small steps across a single LO band. The size of the steps between each tuned frequency depended on the requested redundancy of frequency coverage, which determines the number of frequency settings within the IF of the instrument/spectrometer system. From a data analysis standpoint, the reduction to a SSB spectrum was most reliable when redundancy was high. This had to be balanced with the observing efficiency penalties imposed by the dead times associated with tuning to each new LO frequency. For most sources, a reliable reduction of the line spectrum, at least 5 frequency settings within the IF bandwidth, i.e. a redundancy of 4, was recommended. The spacings between the different LO frequencies contained a small random component to prevent aliasing which could occur in the reduction of the multiple double-sideband measurements to a single sideband (SSB) spectrum in a sideband deconvolution process. The Observer could choose to do a full or a limited range of frequencies in a given LO subband, but only one such subband per observation.

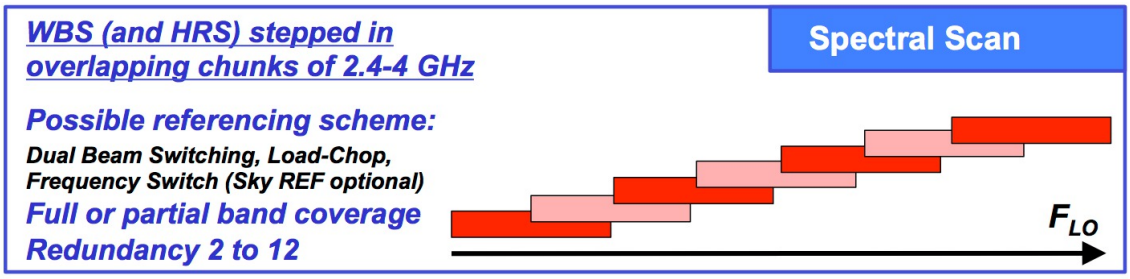

Figure 3.12: Timeline of HIFI Spectral Scan AOT.

The Spectral Scan AOT was offered with Dual Beam Switch, Frequency Switch (with or without sky reference measurements for optimal baseline corrections), and Load Chop (with or without sky reference) calibration modes. Standard Position Switching was not offered, as the efficacy to correct baseline artefacts in total power mode without additional chopped reference measurements (frequency switch or thermal load) was generally low.

Sideband deconvolution is carried out in the HIFI pipeline, and also provided as an interactive tool in HIPE in case the user wished to process the input spectra outside of the pipeline, e.g. additional baseline cleaning to prevent ghost lines or ringing due to unflagged regions of spurious response (see also the HIFI Data Reduction Guide. This latter was also handled in the framework of the Spectral Scan Highly-Processed Data Products (HPDPs) available in the HSA (Section 6.2)

Spectral Scans with Dual Beam Switch (`HifiSScanModeDBS` and `HifiSScanModeFastDBS`)

The use of DBS reference inherited all the advantages and restrictions from the single point dual beam switch observing mode (Mode I-2, Section 3.3.1). The combination of telescope and chopper motions followed the same scheme and timing constraints as this mode. This implied also that this spectral scan mode could only be applied to astronomical sources that were smaller than the chop angle of 3 arc minutes. Each LO frequency was separately calibrated against the internal hot/cold loads of the instrument. Extra observing time (double) was spent at LO frequencies where the system temperature is significantly degraded with respect to the rest of the LO subband being used.

The mode was ideal (in fact optimised) for complex spectra, for which the Observer could select the best frequency redundancy, where higher redundancies increased the fidelity of the SSB spectrum after data reduction. The fast chop and continuum timing options allowed for continuum measurements which were not optimal with the Frequency Switch mode. Thus this mode was also the most useful for measuring absorption lines. The scheme has similar restrictions to other dual beam switch (DBS) measurements. For extended sources the chopped beam could pick up the extended emission or absorption, resulting in chopped line profiles. Fast chop was recommended in bands when the stability (Allan) times at the goal resolution of the system are less than 2 seconds, reducing the observing efficiency. High redundancies were recommended for rich or complex (emission plus absorption, and broad lines), also reducing efficiencies. Under these conditions the Observers may have opted for the Load Chop mode.

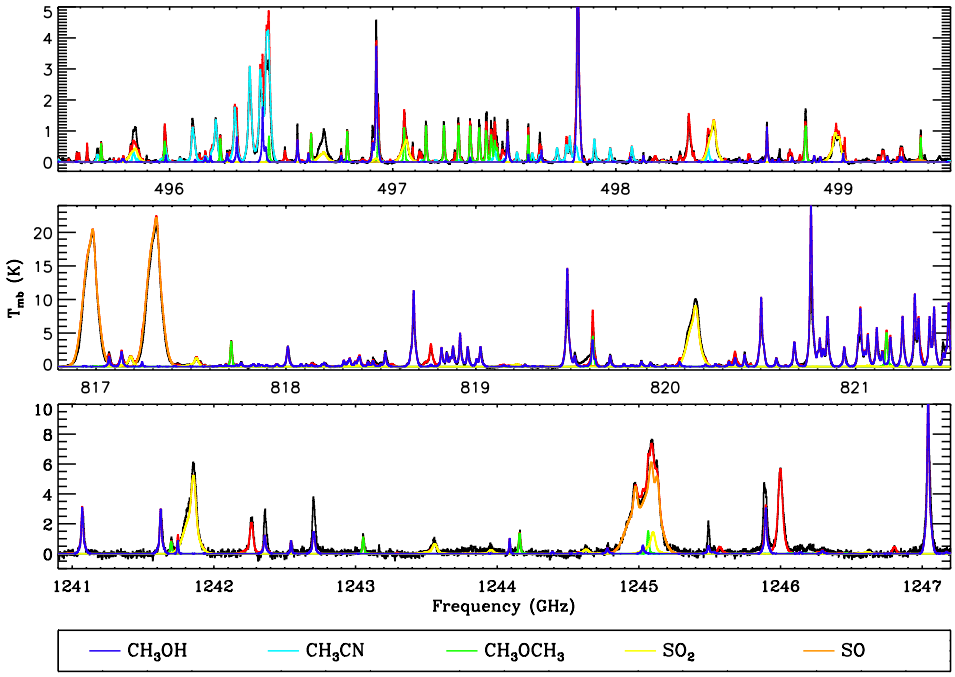

Figure 3.13: A HIFI Spectral Scan AOT example. Three selected spectral regions from the HIFI spectral survey of Orion KL. The data acquired with the Spectral Scan DBS mode are plotted in black and the full band model is overlaid in red. Other colours correspond to individual molecular fits as labeled in the legend. The overlaid individual fits include emission from all detected isotopologues and vibrationally excited states. [Credit: Crockett et al. (2014)]

Spectral Scans with Frequency Switch (`HifiSScanModeFSwitch` and `HifiSScanModeFSwitchNoRef`)

At each LO frequency of the Spectral Scan, the Frequency Switch mode operated in a fashion similar to a Single Point Frequency Switch observation, inheriting similar advantages and disadvantages. The internal reference is from a nearby frequency, using a small step frequency away from each of the main steps that are taken every 0.5 to 1 GHz. The frequency step or throw was selected by the Observer, from the options as given in Table 3.3.1. The initial design of this mode was to group the reference frequency measurement with several "on" frequency measurements during scanning, thereby improving observing efficiency. Performance validation of this frequency grouping strategy showed that system response often varied too much between neighbouring LO frequencies for a general implementation. At any frequency grouping, the mode was not recommended for Spectral Scans in HEB Bands 6 and 7 due to short Allan times and long LO stabilisation times.

Spectral Scans with Load Chop (`HifiSScanModeLoadChop` and `HifiSScanModeLoadChopNoRef`)

Similar to the Single Point Load Chop mode, the reference used at each frequency covered by the Spectral Scan is from the internal cold load. An OFF sky reference up to 2° away from the target could be used to correct standing waves in the baseline, applying the measurements in a double-differencing that remove the effects of light path differences of sky to mixer and sky to internal load. The mode inherits similar advantages and disadvantages to the corresponding Point AOT mode.

3.3.4 Calibration Modes

Users of the HSA will occasionally find HIFI observations with "Calibration" contained in the `PROPOSAL_ID` column of the query output, and in the AOR label. If the user has queried observations with the Standard Data checkbox selected, then all such calibration observations were taken with the standard Observing Modes as they have been described above, and indicated in the `OBSERVING_MODE_ID` column. These observations often contain quite high quality science data. When the user queries with the Standard Data box unchecked, then the query output will contain (among many other things) AOT-like engineering observations as well with similar Observing Mode names. These are the so-called "`Proc`" modes, which are a procedural adaptation of the standard modes and allowed the Calibration Scientist to set certain setup and timing-related parameters more flexibly than allowed (for practical reasons) in HSpot. Such observations are often otherwise identical to the standard modes, using somewhat different sampling in maps for beam calibrations, or different data readout parameters for performance investigations, and so forth. Unlike the standard modes, their observing parameters are not optimised during time estimation in HSpot, but fixed according to need. Nonetheless, these observations often contain good science quality data, but equally often they do not. More information about specfic observations in these two categories is in the HIFI Calibration Observations Release Note (HIFI-ICC 2013).

Chapter 4

HIFI Calibration

4.1 Intensity calibration

4.1.1 Calibration concept

The flux calibration of the HIFI instrument consists of a 3-point calibration scheme. Measurements on two different internal black bodies (called *hot* and *cold loads*) and on an assumed blank sky positions (also called *OFF* position, in constrast with the *ON* target position) are performed within a time-optimised observing sequence. While the internal load measurement allows someone to calibrate out the instrument bandpass (i.e. the instrument spectral response function), the *OFF* position measurement allows to cancel out to first order the instrument response drift. The overall calibration concept is described in detail in Ossenkopf (2003).

The measurements are sequenced assuming an instrument response drift model relying on the Allan variance formalism (Schieder & Kramer 2001, Ossenkopf 2008, see also Section 3.2). From the Allan variance, we can calculate the maximum time between two measurement phases in order to remain in a radiometrically-dominated noise regime, and not be affected significantly by drift noise. This in turns defines the timing of the observing sequence (see also Morris & Ossenkopf 2017).

The application of this calibration scheme in the HIFI pipeline puts the counts measured on a given astronomical source into an antenna temperature scale called T_A^*. This scale is commonly used at radio-telescopes (e.g. Kutner & Ulich 1981, Wilson et al. 2009) and calibrates the source as if it were located next to the instrument internal loads (i.e. it already corrects for the rearward losses, see e.g. Kutner & Ulich 1981). It is fundamental to understand that this scale is HIFI-centric and does not yet take into account the optics between the HIFI focal plane and the sky (in practice, the prime and secondary telescope reflectors). **As such, users of the HIFI data should bear in mind that the T_A^* scale is not readily usable to perform scientific analysis of the HIFI data. Conversion into a more physical scale (typically main beam antenna temperature, $T_{\rm mb}$, or Jansky) that takes into account the optical response and losses of the *Herschel* telescope is needed**. This latter involves the assessment of the telescope coupling efficiencies, which is obtained using a prime calibrator as absolute flux scale reference – for HIFI this prime calibrator is the planet Mars. This is further described in Sections 4.3 and 4.5.

Another aspect fundamental to the HIFI intensity calibration is that HIFI is a double-sideband (DSB) instrument (see Section 2.1.2). As such its calibration has to take into account the respective re-

sponses in the two sidebands in order to accurately calibrate line intensities that belong to the respective spectral ranges. This peculiarity has an important consequence: the calibration concept of HIFI is optimised for *line* calibration. In practice, because the DSB is the sum of two single-sideband spectra, the resulting continuum is approximately twice its real amount, while the line intensities are at their correct levels. Since HIFI is actually able to detect continuum with a good accuracy (see Section 5.8.2), this is important to bear in mind when performing science relying on both line and continuum information (e.g. for absorption studies). Section 5.8 gives recipes on how to deal with such situations.

4.1.2 Calibration parameters

Following the formalism described in Ossenkopf (2003), the calibration equation for signal arising from one particular single sideband (SSB) can be written as:

$$J_{\mathrm{ON,ssb}} - J_{\mathrm{OFF,ssb}} = \frac{\eta_h + \eta_c - 1}{\eta_{\mathrm{sf}}\eta_l G_{\mathrm{ssb}}\omega_{\mathrm{ssb}}} \times \frac{c_{\mathrm{ON}} - c_{\mathrm{OFF}}}{c_h - c_c} \times (J_h - J_c) \tag{4.1}$$

where J is the Rayleigh-Jeans equivalent radiation temperature at a given frequency, and c the counts measured over the respective phases (*ON*, *OFF*, hot and cold).

The calibration parameters entering the equation are:

- the sideband gain ratio G_{ssb}
- the hot and cold load coupling coefficients η_h and η_c
- the source coupling coefficient to the telescope η_{sf}
- the forward beam coupling efficiency η_l
- the optical standing wave ω_{ssb}, expressed as a modulation factor to the gain

Note that HIFI uses a normalised sideband gain ratio formalism, whereby $G_{\mathrm{ssb}} = G_{\mathrm{usb}} = \gamma_{\mathrm{usb}}/(\gamma_{\mathrm{usb}} + \gamma_{\mathrm{lsb}})$ for calibration of lines in the Upper Sideband (γ_{usb} and γ_{usb} being the system gains in the respective sidebands), and $G_{\mathrm{ssb}} = G_{\mathrm{lsb}} = 1 - G_{\mathrm{usb}}$ for calibration of lines in the Lower Sideband. This is explained in more details in Higgins et al. (2014).

Because G_{ssb} is close to 0.5 (i.e. HIFI was for most cases a sideband gain balanced system), one can see in the above equation that the continuum intensity in calibrated spectra will appear as approximatively twice its real value. This is because the continuum information is the superposition of the Upper and Lower sideband fluxes (i.e. a Double Sideband – DSB – flux) while the line is on a Single-Sideband scale – this is equivalent to say that the DSB continuum is about twice the SSB continuum.

In this formalism, the standing wave contribution is considered to be multiplicative, and not additive. This is based on the observational fact that residual (optical) standing waves in the calibrated data enter as a gain component. This has mostly effects on the intensity calibration uncertainty computation (see Section 5.8.2).

Finally, as explained in the previous sub-section, the HIFI pipeline cannot calculate the source coupling efficiency η_{sf}, and so like for other radio-telescopes, the data delivered in the top level products is on the T_A^* scale.

Table 4.1: Overall HIFI Frequency Accuracy Budget (Pre-Launch)

Assumptions								
Frequency (GHz)	480	640	800	960	1120	1250	1410	1910
LO syst. freq. acc. (MHz)	0.024	0.032	0.04	0.048	0.056	0.0625	0.0705	0.0955
IF up-convert freq. acc. (MHz)							0.050	0.050
WBS								
AOS freq. cal. acc. (MHz)	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
Instr. freq. acc. (MHz)	0.12	0.13	0.14	0.15	0.16	0.16	0.22	0.25
HRS Normal								
HRS freq. acc. (kHz)	5	5	5	5	5	5	5	5
Instr. freq. acc. (kHz)	29	37	45	53	61	68	126	151

4.2 Spectral Calibration

The (super)heterodyne principle, mixing an incoming signal with a local oscillator (LO) to down-convert the signal to an intermediate frequency band, allows detectors to be built for a conveniently low frequency range. High frequency resolution can be achieved, and the accuracy depends on the stability of the spectrometers and of the LO. In the case of digital spectrometers in which the electronics clock is locked to the LO, such as HIFI's HRS, the accuracy depended mainly on the LO stability. The WBS acousto-optical spectrometers had extra calibration steps for the CCDs and this was the main contribution to WBS frequency calibration uncertainty.

The LO was of fundamental importance to HIFI. Mis-behaviour, such as harmonic generation, could introduce data artefacts; see section 5.3.

It should be clear that frequency accuracy and resolution are separate system properties. The HIFI instrumental response spred an essentially unresolved spectral line over a number of channels, which was the resolution. By accuracy we refer to the channel labelling: how close was the centroid of the frequency PSF to the true frequency of the line?

The pre-launch accuracy budget for frequency calibration (table 4.1), with technical descriptions, is in Herpin & Teyssier (2003). These goals were achieved, fortunately, and section 5.7 discusses in-flight tests done to confirm the frequency calibration accuracy.

4.2.1 Master Oscillator and Local Oscillator

The HIFI master oscillator was a temperature-controlled 10 MHz crystal oscillator (OCXO). Although the temperature-control power circuit developed a fault during the mission, the nominal operating conditions of the oscillator were not exceeded. The nominal purity of the OCXO was:

$$\delta\nu/\nu \approx 10^{-8}$$

This fundamental frequency was upconverted with phase-locked synthesizers and further up-multiplied to produce the LO signal used to down-convert the sky signal. Although multipliers used in LO generation introduced a negligible frequency error, the up-converter oscillator in the Local Oscillator Source Unit (LSU) did have a small spectral width of 30 kHz–100 kHz depending on observing

band (Herpin & Teyssier 2003). The HEB bands used an additional free-running 10.4 GHz local oscillator with expected stability of 3 to 4 $\times 10^{-6}$, therefore adding 30 kHz of uncertainty in bands 6 and 7.

4.2.2 High Resolution Spectrometer

After down-conversion, the signals were in an intermediate frequency (IF) range of 4–8 GHz. Detection took place in this band.

The HRS was a digital auto-correlating spectrometer (ACS). Labelling the frequency channels in the output of an ACS was a matter of knowing how many samples went in to the correlation and the time between them. The HRS electronics were clocked in sync with the fundamental 10 MHz oscillator, hence the IF frequencies assigned to the output detector channels depended on the accuracy of the OCXO and the LO, which was $\lesssim$ 50 kHz. The spectrometer resolution (spectral PSF) depended on observing mode and was known, see table 2.3.

4.2.3 Wide Band Spectrometer

The WBS acousto-optical spectrometers detected the output of the Bragg cells on four linear CCD arrays (per polarisation). IF frequency increased nonlinearly along the CCDs. Frequency calibration involved determining the IF frequency of each CCD pixel. An internal COMB spectrum, consisting of narrow tones separated by 100 MHz and locked to the OCXO, was generated and detected on the CCDs. During processing by the HIFI pipeline, the COMB peaks were fitted and used to derive a pixel-to-IF-frequency polynomial mapping. The WBS resolution was also computed from the unresolved COMB lines; a typical value was 1.1 MHz. The combined uncertainties of fitting to the COMB peaks and then interpolating the frequency mapping over the IF dominated the total frequency calibration uncertainty for WBS. In-flight end-to-end tests (5.7) demonstrated a total uncertainty of $\pm$ 100 kHz for the WBS frequency calibration.

4.3 Beam Calibration

HIFI is a sub-millimeter receiver mounted on a telescope that was designed for operation at infrared wavelengths. Finding the optimum sub-millimeter wave illumination for such large telescopes is simplified by their effectively perfect optical surfaces and long focal lengths, but is complicated by telescope optics which have been optimised for non-tapered beams. So there is a need to understand the effects of blockage and the tradeoffs with edge taper, the latter defines the illumination of the telescope. A low taper will illuminate the telescope more evenly, but introduces more spillover effects which lead to stronger side lobes in the antenna pattern; a high taper reduces the side lobes at the expense of a broader beam, one is loosing spatial resolution as one is no longer making optimal use of the outermost region of the telescope mirror.

The total forward beam efficiency is the forward coupling efficiency to a source which fills the entire beam including the side lobes and will be determined by overall phase errors, spillover effects, blockage, cross-polarisation, random phase errors across the reflector surface, and ohmic loss.

The terms which will depend on telescope illumination the most are spillover and blockage. Spillover is power passing beyond the maximum telescope radius, while blockage is power blocked by the

Table 4.2: Geometrical parameters of the Herschel telescope. Column 3 gives the purely geometrical blockage relative to the effective aperture. The total geometrical blockage is the sum of M1 central hole (2.9%), 3 Hexapod holes (0.5%), shadow of hexapod legs + M2 barrel (1.7%), shadow of hexapod legs (path from secondary to primary) (2.5%). The effective blockage will still be larger than the total geometrical blockage (from Kramer 2008).

	diameter	geometrical blockage
Primary:	3.5 m	
Effective aperture	3.28 m	
Secondary:	308.1 mm	0.9%
Central hole in primary:	560 mm	3%
Total blockage		7.7%

central obstruction. Central blockage may come either from the secondary itself or, more commonly, from the Cassegrain hole or mirror support in the centre of the primary. For centrally blocked optics and tapered beams the spillover and blockage product efficiency peaks at some intermediate edge taper. If the edge taper is too large a disproportionate amount of power is blocked by the central obscuration; if the edge taper is too small a disproportionate amount of power spills off past the edge of the telescope.

The Herschel telescope is an axi-symmetric 3.5 m diameter Cassegrain telescope consisting of a parabolic primary (M1) and hyperbolic secondary (M2). There is no mechanism for refocussing once the telescope is in orbit. The operating wavelength ranges from 80 µm to 670 µm. A few important parameters are listed in Table 4.2 (Kramer 2008). From this table we get $D = 3.5\,\mathrm{m} > 2a = 3.28\,\mathrm{m}$ and $2b = 0.56\,\mathrm{m}$. The blockage factor is given by $f_b = \frac{b}{a}$. The total purely geometrical obscuration ratio was specified to be less than 8% (Jellema 2015).

4.4 Spatial calibration revisited

We here revisit the spatial calibration of HIFI, previously derived by Kramer (2008) assuming a Gaussian beam shape. We drop the latter assumption and derive corresponding expressions for a general beam shape. In the limit of a Gaussian beam, all derived expressions converge to those given by Kramer (2008). See Jellema (2015) for further detail.

Denoting the beam pattern as $P(\phi, \theta)$ (normalised to unity at nominal pointing, i.e., at 0,0) and the flux distribution of the source, expressed in Rayleigh-Jeans brightness temperature, at wavelength λ as $T(\phi, \theta, \lambda)$, the antenna temperature on the T_A^* scale (4.1.2 observed at position (ϕ_0, θ_0) is the convolution of the beam pattern with the source structure:

$$T_A^*(\phi_0, \theta_0) = \frac{\iint_{4\pi} P(\phi - \phi_0, \theta - \theta_0)\, T(\phi, \theta, \lambda)\, \mathrm{d}\Omega}{\iint_{4\pi} P(\phi, \theta)\, \mathrm{d}\Omega} = \frac{P \otimes T}{\eta_l \Omega_a} \tag{4.2}$$

where the integrals are performed over the whole sky (solid angle 4π). For a homogeneous source that completely fills the beam, T_A^* equals the brightness temperature of the source.

The denominator in (4.2), Ω_a, is referred to as the antenna solid angle, which together with the observing wavelength defines the effective area of the telescope, A_e, a relation called the "antenna theorem":

$$A_e \Omega_a = \lambda^2. \tag{4.3}$$

In the case of a telescope such as *Herschel*, in which the entrance pupil has a well-defined geometric area A_g[1], the aperture efficiency η_A can be defined as

$$\eta_A = \frac{A_e}{A_g}. \tag{4.4}$$

It then follows:

$$T_A^* (\phi_0, \theta_0) = \frac{\eta_A}{\eta_l} \frac{A_g}{\lambda^2} \iint_{source} P (\phi - \phi_0, \theta - \theta_0) \, T (\phi, \theta, \lambda) \, \mathrm{d}\Omega. \tag{4.5}$$

The antenna temperature relates to the flux density $S_{\nu,tot}$, i.e. the power radiated per unit area and per unit frequency from a radio source at a given frequency, via

$$T_A^* (\phi_0, \theta_0) = \frac{\eta_A}{\eta_l} \frac{S_{\nu,tot} A_g}{2k} \frac{\iint_{source} P (\phi - \phi_0, \theta - \theta_0) \, T (\phi, \theta, \lambda) \, \mathrm{d}\Omega}{\iint_{source} T (\phi, \theta, \lambda) \, \mathrm{d}\Omega} = \frac{\eta_A}{\eta_l} \frac{S_{\nu,tot} A_g}{2k} K. \tag{4.6}$$

with a conversion factor from Kelvin to Jansky of $2k/(\eta_A A_g)$ and the flux dilution factor K.

4.4.1 Aperture Efficiency

The aperture efficiency η_A defined in 4.4 is the forward coupling efficiency to a point source in the far field. The peak aperture efficiency drops with increasing central obscuration because power from the feed does not reach the primary.

The aperture efficiency can be determined via the observed peak antenna temperature T_A^* of a point-like source when its total flux density $S_{\nu,tot}$ is known. If we express the aperture efficiency as being the ratio of geometrical to effective area of the aperture, we have

$$\eta_A = \frac{A_e}{A_g} \quad \text{and} \quad \frac{\eta_A}{\eta_l} = \frac{2k}{A_g} \frac{T_A^*}{S_{\nu,tot} K} \equiv \frac{2k}{A_g} \frac{1}{\chi_{\mathrm{PSS}}} \tag{4.7}$$

thereby also defining the point source sensitivity in Jansky per Kelvin which is often used instead of the aperture efficiency: $\chi_{\mathrm{PSS}} = S_{\nu,beam}/T_A^*$ (Kramer 2008).

[1] This is simply the geometric area of a circle, but using the effective diameter of 3.28 m of *Herschel*.

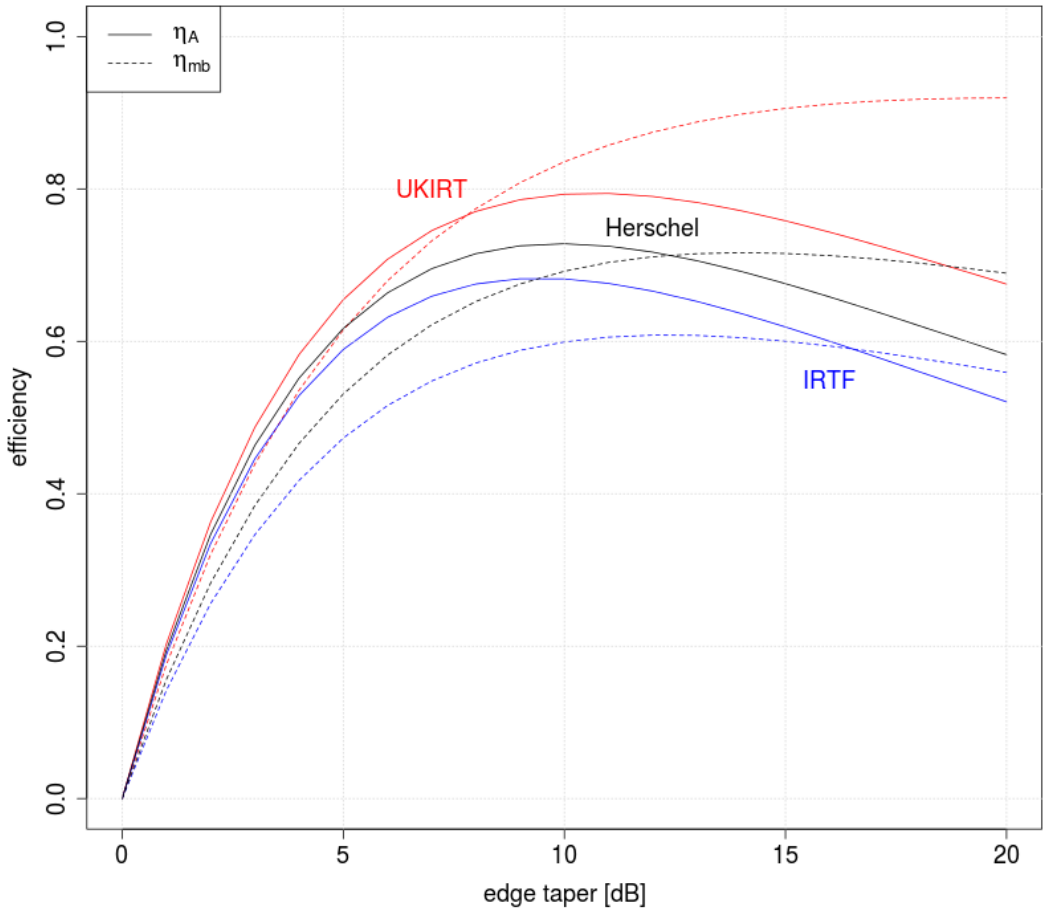

Figure 4.1: Expected variation of main beam (dashed) and aperture efficiency (solid) curves for the IRTF, UKIRT, and Herschel Telescope as function of edge taper.

4.4.2 Main Beam Efficiency

On the telescope, the usual approximation is that the beam on the sky has a Gaussian power distribution with additional far-out side lobes. A correction factor for the beam coupling to sources is commonly called the "main beam efficiency", η_{mb}, which is the forward coupling efficiency to a uniform disk filling the main beam[2], including the geometrical solid angle beam efficiency, telescope losses, and deviations from an ideal Gaussian beam. It is typically measured by comparing the apparent and known temperatures of a planet or the Moon. As the Moon was not available as an astronomical target, all measurements of the HIFI beam efficiencies were based on observations of Mars. The actual main beam shape is not exactly Gaussian, and may have close-in side lobes and stray beams, so the main beam efficiency depends on the calibration source size and temperature distribution. The model by R. Moreno (http://www.lesia.obspm.fr/perso/emmanuel-lellouch/mars/) was used to calculate expected brightness temperatures of Mars in the HIFI frequency bands for given observing dates.

Since interesting small scale structure has a size between a point source and the main beam size,

[2]Defined as the central part of the up to the first nulls, i.e. excluding side lobes.

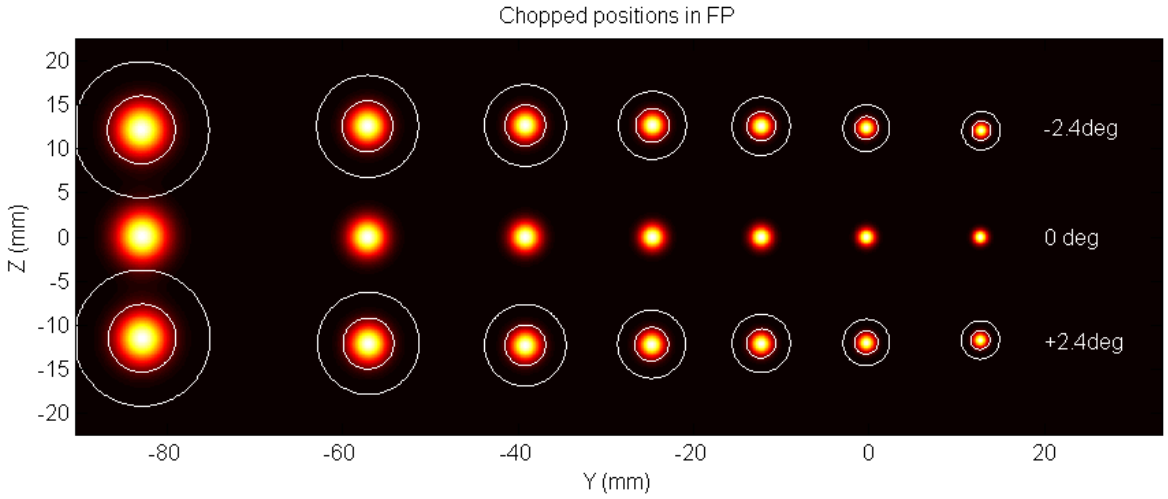

Figure 4.2: Chopped positions in the focal plane. Band 1 is at the left side of the diagram, band 7 at the right edge.

a useful general taper lies between the optima given by the aperture efficiency η_A and by the main beam efficiency η_{mb}. Curves of ideal η_A and η_{mb} for a few infrared telescopes with slightly different diameters and blocking factors are plotted in Fig. 4.1 following Harris (1988). The designed edge taper for Herschel is 10.9 dB (Jellema 2015).

4.4.3 HIFI focal plane geometry

The nominal focal plane geometry of the HIFI optics is depicted in Fig 4.2. 1. HIFI has 7 mixer bands spatially separated along the spacecraft Y-axis. Each mixer band contains two mixers operating on orthogonal polarisations, referred to as H and V designating Horizontal respectively. Vertical mixer polarisation, which are nominally co-aligned in the focal plane. The HIFI Focal Plane Chopper provides focal-plane chopping functionality for each band around its central position with throw along the spacecraft Z-axis. The two chopped positions which are separated by $1.5'$ at the sky are indicated in Fig 4.2 (Jellema et al. 2008).

The measurements of the focal plane geometry was mainly aimed at deriving the HIFI component of the Spacecraft-Instrument Alignment Matrix (SIAM), which defined the position of the HIFI apertures with respect to the Spacecraft boresight (Section 5.9). The structure of the HIFI SIAM entries is explained in Appendix A.

4.5 Calibration Sources and Models

The prime calibrator for HIFI is the planet Mars. As seen from *Herschel*, it had an apparent size between 5 and 9 arcsec depending on the observing date.

The continuum radiation emitted by the surface of Mars is not fully constant. Dust storms and rotational absorption lines, caused by CO and H_2O in Mars tenuous atmosphere, disturb this contin-

uum radiation. Although this is usually treated in canonical Mars models, this does not have an important impact in the accuracy of the calibration. The Mars continuum model used for the absolute flux calibration of HIFI was provided by Raphael Moreno, and is based on the work of Rudy et al. (1987). The initial model was developed by E. Lellouch and H. Amri, and it is available in its version 1.2 to the community at the following link: http://www.lesia.obspm.fr/perso/emmanuel-lellouch/mars/

The program computes the Mars brightness temperature and flux for the specific dates and frequencies of the Martian observations. The output parameters for each frequency are the HPBW, the filling factor f of Mars in the beam, the total flux (Jy) emitted by Mars, the associated mean Planck brightness temperature over the planet ($T_{\rm b}$), the flux in the main beam (in Jy), the associated Planck brightness temperature ($T_{\rm b,beam}$), and the main-beam Rayleigh-Jeans temperature ($T_{\rm mb}$). This latter is the most useful parameter for calibration purposes, since it can be directly compared to observations.

The code uses surface and subsurface temperatures taken from the European Martian General Circulation Model (http://www.lmd.jussieu.fr/mars.html) and Martian ephemerides from Institut de Mécanique Céleste et de Calcul des Ephémérides (IMCCE) (http://www.imcce.fr). The disk is split on a 50 × 50 grid, each point of the grid having its own latitude, longitude, and local time. Brightness temperatures and local fluxes are then calculated and convolved with the gaussian beam of interest.

Further detail about the Mars model values used at the particular dates of the HIFI observations can be found on the *Herschel* Legacy Calibrator page.

4.6 Calibration Uncertainty Model

The Intensity Calibration uncertainty computation is based on the general intensity calibration equation introduced in 4.1. Each parameter entering the equation introduces a certain error component which is then propagated taking into account the respective random or systematic nature of the errors. Following the formalism in Ossenkopf (2015), all components but the planetary model uncertainty can be treated as random errors and therefore be considered statistically independent and added quadratically. The parameters entering the uncertainty models are:

- the mixer sideband gain ratio
- the internal load coupling efficiencies
- the internal load temperature sensor readouts
- the optical standing wave contribution coming respectively from the paths between the mixer horns and the internal loads, and between the mixer horns and the diplexer roof-top mirrors (this layer only applicable to bands 3, 4, 6 and 7)

While the uncertainties on the internal load coupling efficiencies and temperature sensor readouts are based on measurements performed pre-launch (Larsson & Teyssier 2008), the contribution from residual optical standing wave modulations has been assessed after the mission completion throughout an extensive data mining exercise (see also Teyssier 2017c and Rodríguez-Coira & Teyssier 2015).

Finally, the mixer sideband gain ratio has been estimated based on a combination of pre-launch measurements (Higgins 2011, Higgins et al. 2014) and of in-flight spectral scan data (Kester et al. 2016) – see Section 5.8.1 for further details.

The detailed equations describing how each of the calibration parameters enters the total uncertainty can be found in Ossenkopf (2015), and are not reproduced here. Section 5.8.2 indicates how this model translates into intensity calibration uncertainties for each band as a function of LO frequency.

Chapter 5

HIFI in-flight performance

5.1 Sensitivity

Like all heterodyne detectors, the HIFI sensitivity is measured in terms of *System Noise Temperature*, or $T_{\rm sys}$, in units of Kelvin. This temperature is assessed through the contrast of counts measured on the respective hot and cold loads (so-called Y-factor) – higher ratio between those two indicates lower $T_{\rm sys}$ and therefore better sensitivity (see e.g. Ossenkopf 2003).

Figure 5.1 summarises the $T_{\rm sys}$ measured as function of the LO frequency. As can be seen, the $T_{\rm sys}$ is usually a smooth function of frequency, although at several places it experiences sudden increases. These ranges correspond to frequency ranges that could not optimally be tuned due to shortage of LO output power, and therefore turned out to offer relatively poor sensitivity compared to the baseline. While the $T_{\rm sys}$ in SIS bands (1 to 5) is close to monodically increasing with frequency (as a consequence of the quantum noise limit applying to each mixer junction technology), it is fairly flat and homogeneous for HEB bands (6 and 7). Close-ups of the $T_{\rm sys}$ variation with LO frequencies are also shown in Figures 5.2 to 5.8.

The sensitivity was observed to be highly reproducible over the mission. They are, however, some exceptions in frequency tuning areas where the LO output power was just marginal and so from one tuning to another the automatic power levelling algorithm (see the tuning scheme section of Morris & Ossenkopf 2017) did not necessarily succeed. In case of a failed, or partially failed LO power tuning, certain House-Keeping parameters will take anomalous values, and this fact is used to raise dedicated quality flags and warn the users. The corresponding flag reads "FPU Check:The level of mixer current is Out Of Limit" (see also 6.3.2). Its presence does not imply per se that the data noise will be larger than expected, but data with poorer noise performance will always contain this flag. One noticeable problematic case is that of the HEB bands where the ability of optimally tune can change on very short frequency range. We observed this issue mostly in the following ranges:

- in the LO frequency range 1440–1441.5 and 1447.2–1448 GHz (band 6a)
- in the LO frequency range 1756–1758 GHz (band 7a), it was often the case that the V polarisation at least had trouble to be pumped at all
- in observations targeting the [CII] line at LO settings around 1897 GHz (band 7b). In the early phase of the mission there was frequent problem to pump the mixer in this range, but this got

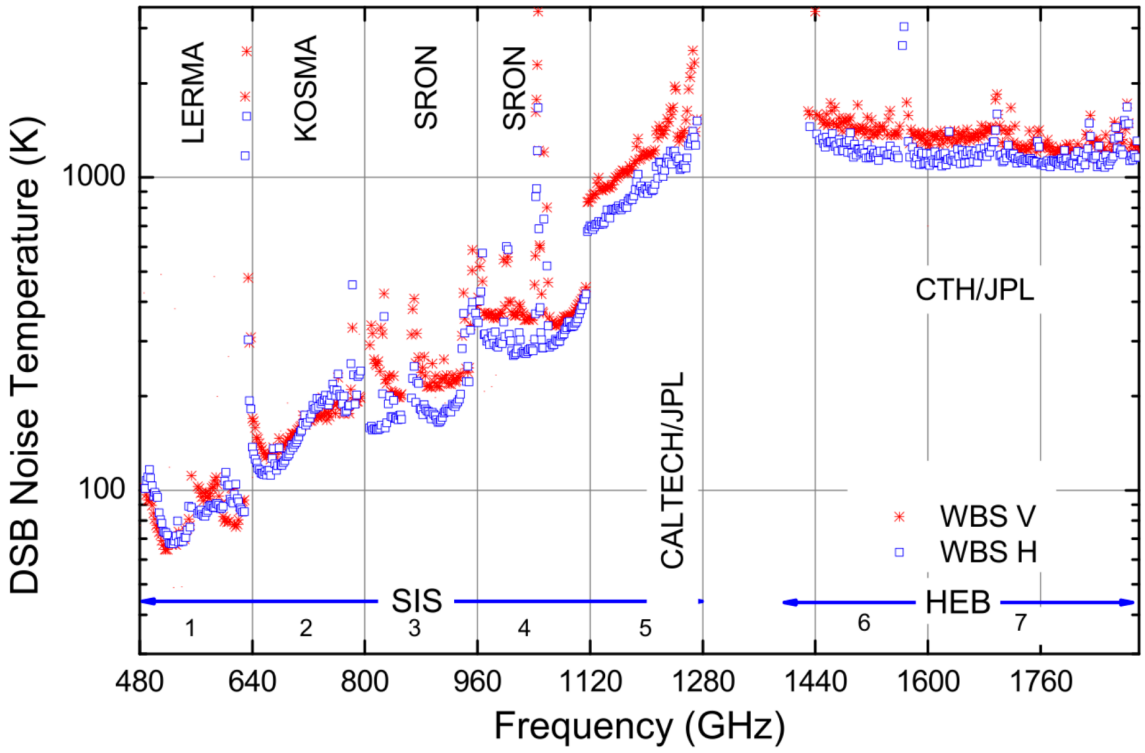

Figure 5.1: HIFI System Noise Temperature as a function of LO Frequency, for the H and V polarisation mixers. Also indicated are the institutes who were responsible for the design and manufacture of each band (from Roelfsema et al. 2012).

fixed from OD 305 onwards. After this date, tunings of the [CII] line have sporadically failed when the tuned frequency fell into the narrow range 1897 ± 0.25 GHz.

The sensitivity is not only a function of LO frequency, but also a function of the Intermediate Frequency. In bands using a beam-splitter injection for the LO signal (see Section 2.3.2), the $T_{\rm sys}$ spectrum is relatively flat, however in bands using diplexers (bands 3, 4, 6 and 7), the $T_{\rm sys}$ has a parabolic shape with a degradation of the sensitivity of order 30% at the edges. In bands 6 and 7, this degradation is larger (factor of 2 difference between the minimum and maximum $T_{\rm sys}$ over the IF), and the best sensitivity is not seen in the middle of the IF but rather at its lower end. Figure 5.9 illustrates some typical $T_{\rm sys}$ spectral shapes in representative mixer bands. The $T_{\rm sys}$ spectrum measured for each observation and tuning is available in the data products (see Section 6.1). A database of $T_{\rm sys}$ spectra sampled at 1 GHz step is also readily available as Ancillary Data Products, see:

http://archives.esac.esa.int/hsa/legacy/ADP/HIFI/HIFI_tsys_data_products/

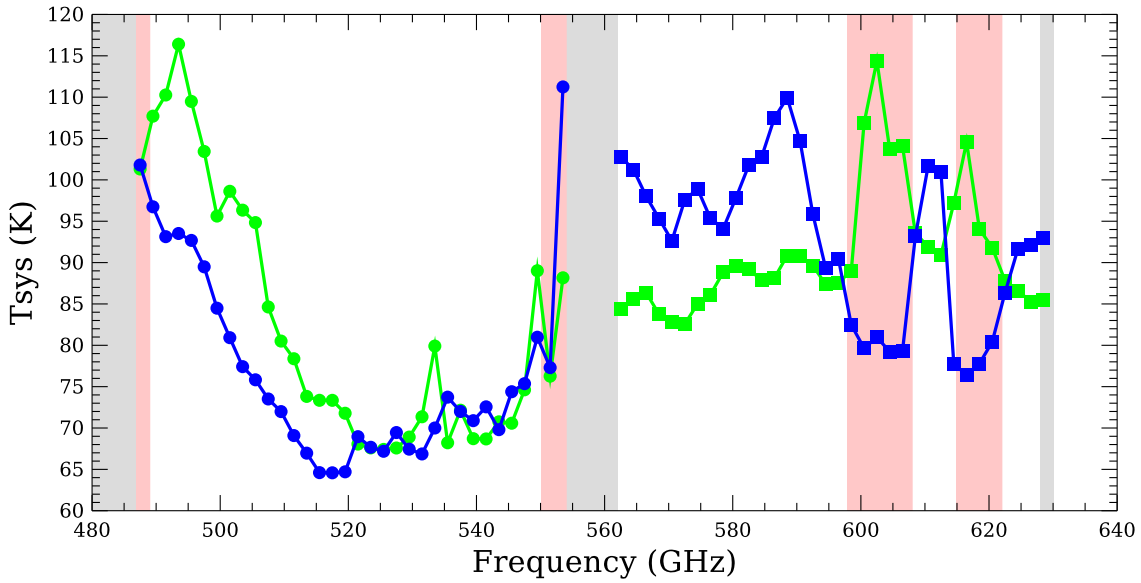

Figure 5.2: HIFI System Noise Temperature as a function of LO Frequency for mixer band 1, for the H (green) and V (blue) polarisation mixers. The red areas indicate ranges of limited LO output power, that could lead to under-pumped or marginally pumped mixers.

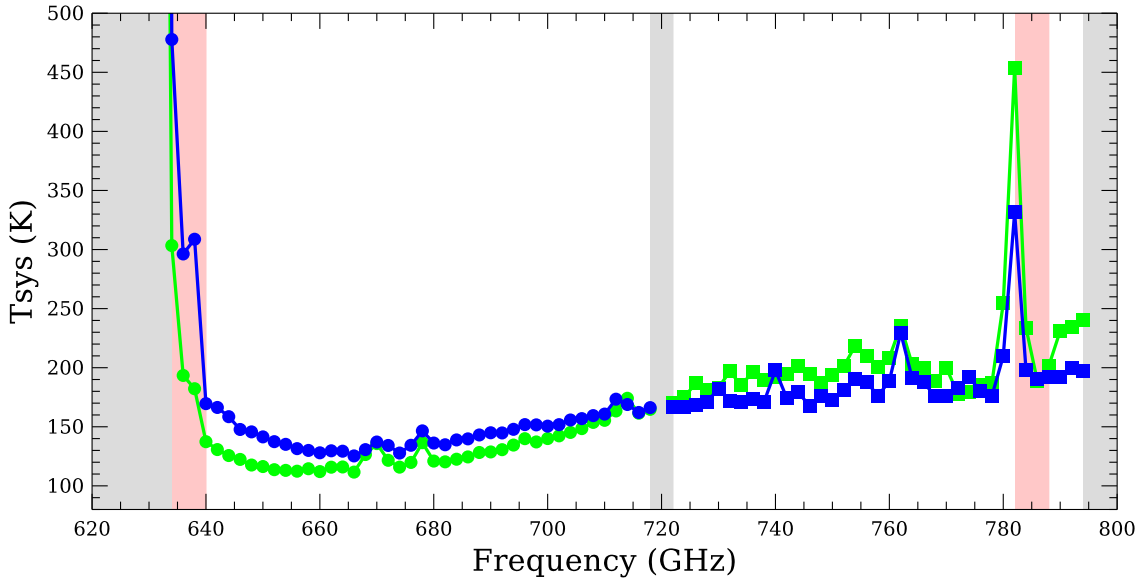

Figure 5.3: Same as Figure 5.2 for Band 2.

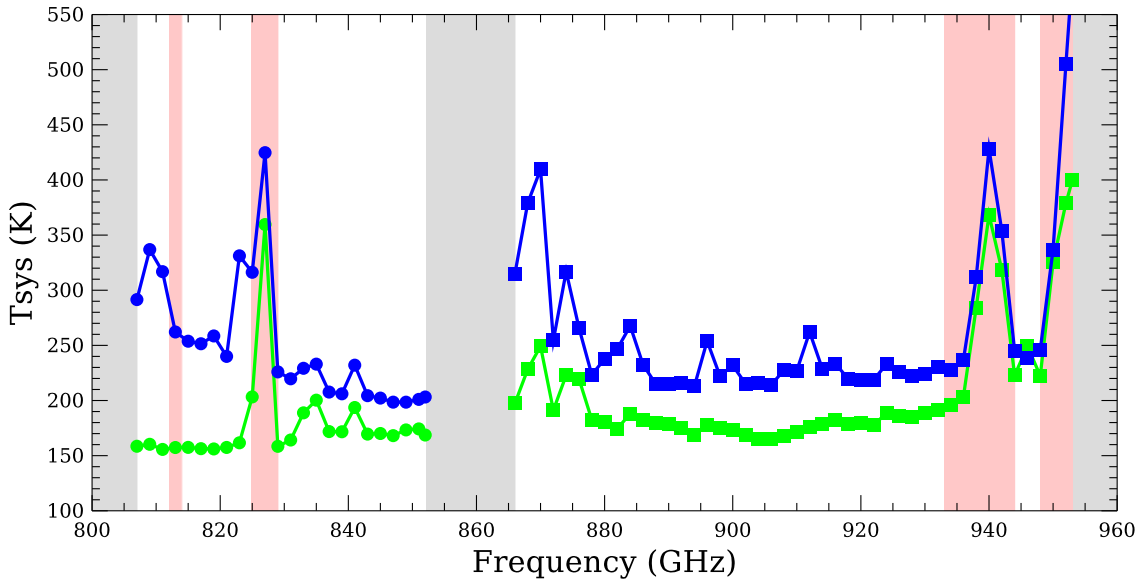

Figure 5.4: Same as Figure 5.2 for Band 3.

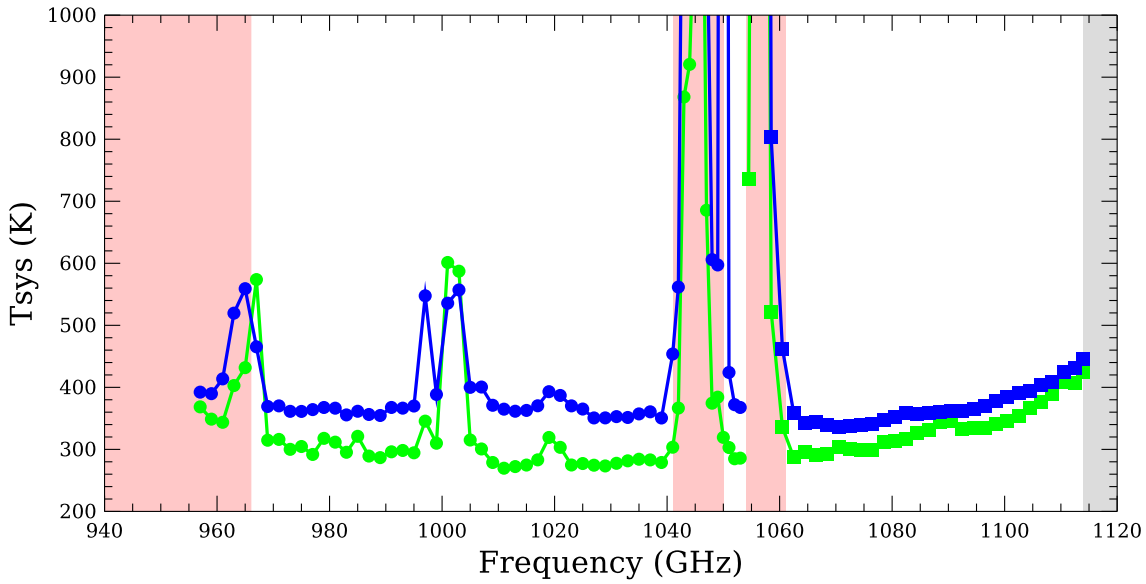

Figure 5.5: Same as Figure 5.2 for Band 4.

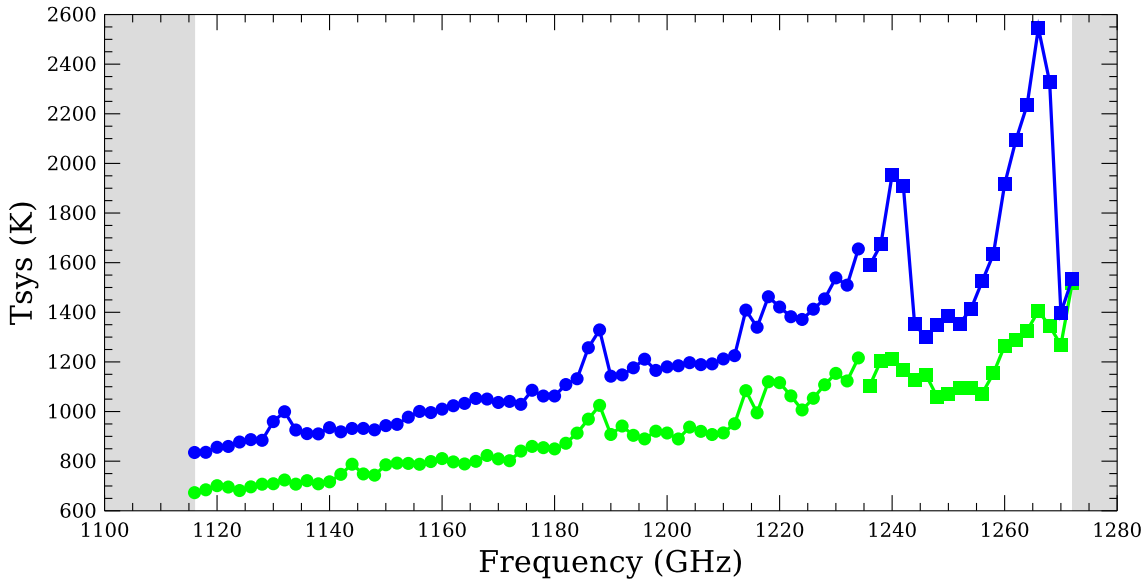

Figure 5.6: Same as Figure 5.2 for Band 5.

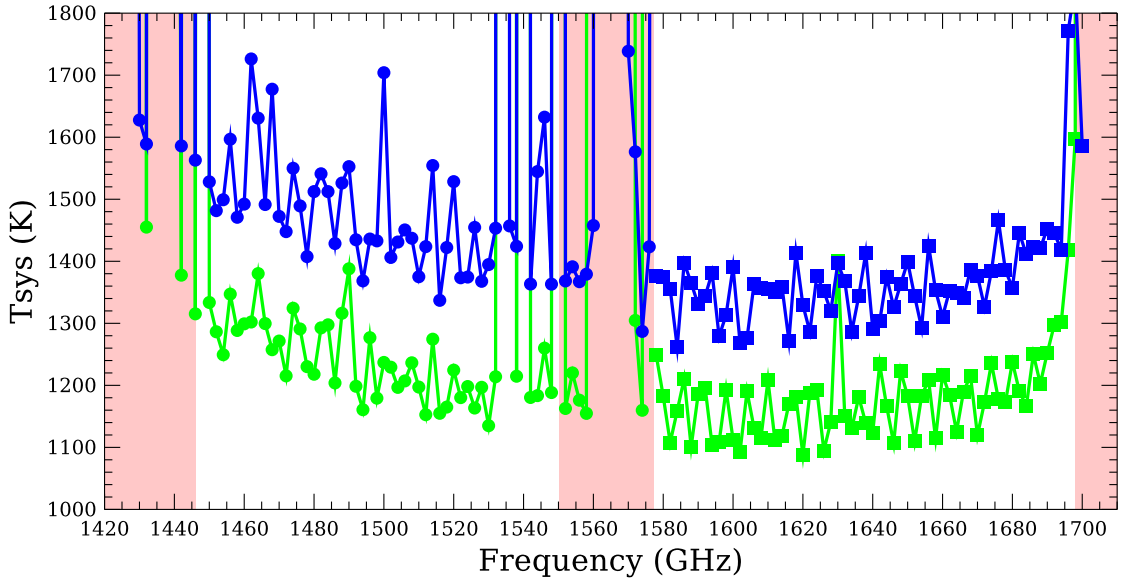

Figure 5.7: Same as Figure 5.2 for Band 6.

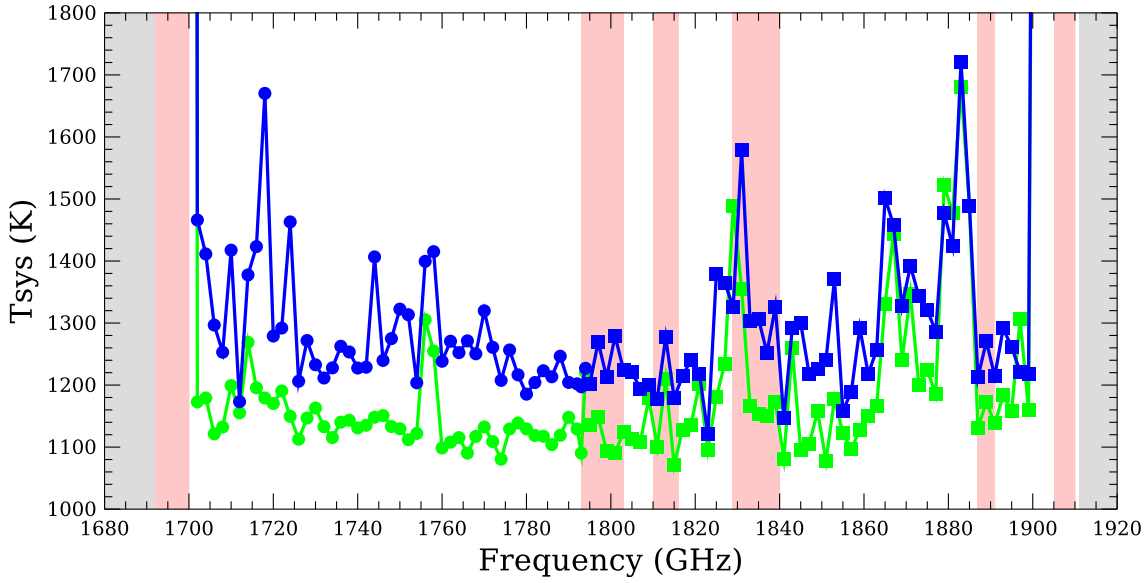

Figure 5.8: Same as Figure 5.2 for Band 7.

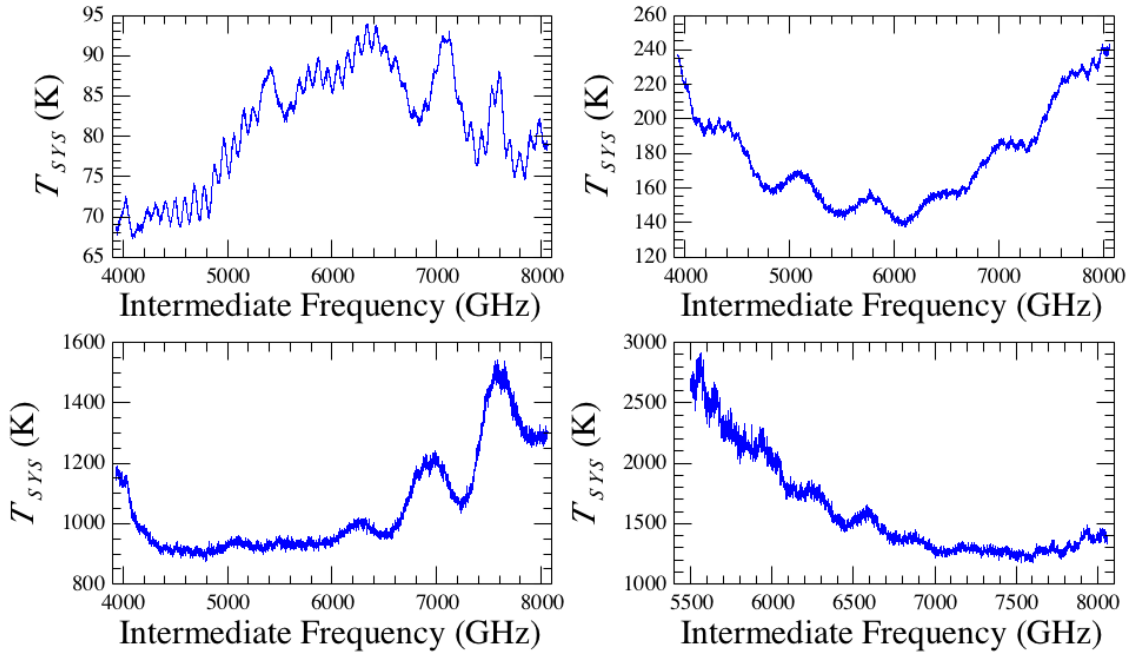

Figure 5.9: Representative HIFI System Noise Temperature spectra as a function of the IF, in bands 1 (upper left), 3 (upper right), 5 (lower left) and 7 (lower right).

Sensitivity in the High-Resolution Spectrometer

The noise achieved in the HRS spectra suffers a slight degradation compared to that of an ideal radiometer at the applicable frequency resolution because of the so-called efficiency of the auto-correlator. For the HIFI high-resolution spectrometers, this efficiency is 0.81. However the efficiency is based on a certain assumption in the signal input level regime, corresponding to signal levels close to those observed on a blank sky. Because the spectrometer level tuning is performed on the internal hot load (100 K), the efficiency will slightly depart from its theoretical value, by an amount related to the Y-factor (ratio of the IF power towards the hot and cold loads respectively). It is estimated that the efficiency could drop as low as 0.75 in bands 1 and 2, with a noise degradation as high as 17% compared to the noise levels predicted by HSpot. The effect is much less in the high-frequency bands due to their low Y-factor.

5.2 Stability Times

As indicated in Section 4.1, the observing sequence of the various measurement phases is timed in order to minimise the drift noise in the calibrated data. Differential measurements are combined to beat this drift and the time elapsed between two phases is driven by the Allan variance time (Ossenkopf 2008) – this time is a measure of the instrument stability. Allan times vary between bands and frequencies within bands. They can also be different depending on the optical path involved in the measurement phases.

Allan variance times have been measured both in the laboratory and in orbit – in general the in-flight performance is better than the pre-flight one, mostly due to the difficulty to reproduce the thermally stable environment offered at the Lagrange 2 location in the laboratory. Table 5.1 gives a summary of the typical Allan variance times offered by HIFI. A full description of the measurements and results can be found in Kooi & V. (2009b) and Kooi & V. (2009a). As can be seen, the instrument stability is usually better in SIS bands than HEB bands, due to the broader band nature of the latter. Bands 3a and 4a are the poorest SIS bands, while band 6b is the best HEB band. Also, the total power mode stability performance is worse than spectroscopic and differential stability performance, as is expected for broadband stability (Ossenkopf 2008) – as such total power stability represent the instrument stability performance for continuum measurement. Differential stability is the best of all, simply because it does not let the drift contribute on short time scale. Frequency-Switching offers the poorest differential stability, mostly due to the short-term variation of instrument bandpass when moving from one frequency phase to another (see Chapter 3).

The above Allan variance times are representative of an instrument that has thermally stabilised, i.e. is not subject to short-term drifts linked to electronic switches. As such they are not fully representative of the instrument state immediately after a Local Oscillator band switch (we recall that HIFI could be operated with only one LO band at a time). Because the thermal drift of the LO electronic chain is highest immediately after a band switch, a dead-time had to be introduced prior to any scientific observation. This dead-time varied with the band, and whether the science observations to follow the switch-on was to be using a slow- or a fast-referencing scheme (see Chapter 3). For a slow-referencing scheme, the switch-on dead-time was of order 30 minutes in SIS bands and 50 minutes for HEB bands. For a fast-referencing scheme, it was of order 4-5 minutes in SIS bands and 15 minutes for HEB bands. This effect can also occur when changing LO frequency within a given band, even well after the band switch-on. This can result in degraded stability on short time scale and is illustrated in 5.3.1.

Table 5.1: HIFI in-orbit stability performance summary (from Kooi & V. 2009b and Kooi & V. 2009a). The numbers correspond to the typical Allan variance time in seconds measured over the whole IF bandwidth of the respective Wide-Band Spectrometers. Differential stability is measured for Double Beam Switching ([1]), Frequency Switching ([2]) and Load Chop ([3]). "–" indicates that no meaningful result could be derived.

Band	Total power		Spectro		Differential							
					Internal loads		DBS[1]		FSW[2]		LC[3]	
	H	V	H	V	H	V	H	V	H	V	H	V
1a	7	7	78	66	461	527	1800	–	1800	1282	1800	1800
1b	14	11	54	57	778	277	1800	1800	1800	1800	1800	1547
2a	14	21	95	72	706	912	–	1800	1800	1683	1800	1800
2b	4	5	85	86	330	573	1800	1800	1271	1150	1644	–
3a	6	9	47	36	1800	503	–	1800	1334	1075	–	1800
3b	12	14	67	77	–	850	1800	1800	1800	1800	1800	1800
4a	11	7	23	15	–	368	487	1649	372	997	634	1085
4b	9	12	61	63	1800	888	1800	1800	377	1800	1800	1800
5a	19	20	133	94	1800	1631	1800	1800	1584	1600	1800	1800
5b	17	17	94	93	1271	1377	1800	1800	1164	–	1688	1800
6a	7	7	23	25	–	–	1800	1598	1331	1800	1800	1800
6b	5	4	66	61	600	971	1800	1800	1027	983	1666	1800
7a	3	5	58	80	593	455	1538	1800	1461	916	–	1250
7b	7	7	17	20	1113	1066	765	522	1800	–	1800	256

Another effect was observed early on in the mission, whereby the thermal stability of the WBS-V optical unit was affected by the switch of the SPIRE Data Control Unit when HIFI was observing immediately after SPIRE (see Dieleman 2009 and Kooi 2009). The affect was such that the WBS optical unit thermal specification were way above the acceptable limit, so that for the remainder of the mission, it was agreed to never schedule SPIRE before HIFI. As such this effect did not apply to the data taken in the Routine Phase.

5.3 Data artefacts

For the large majority of the data taken by HIFI, the calibration scheme did a very good job in calibrating out the instrumental response signatures. In some cases, however, the data still suffer from residual artefacts in the top level data products (see also Section 6.3). These will take two main forms:

- baseline distortion, resulting in non-flat baseline with localised bumps and ripples arising from residual standing waves
- spurious responses and spurious signals, manifesting as ghost spectral features ("spurs") and that sometimes go accompanied with erroneous intensity calibration

The following subsections give an overview of the zoology of those artefacts, with representative examples.

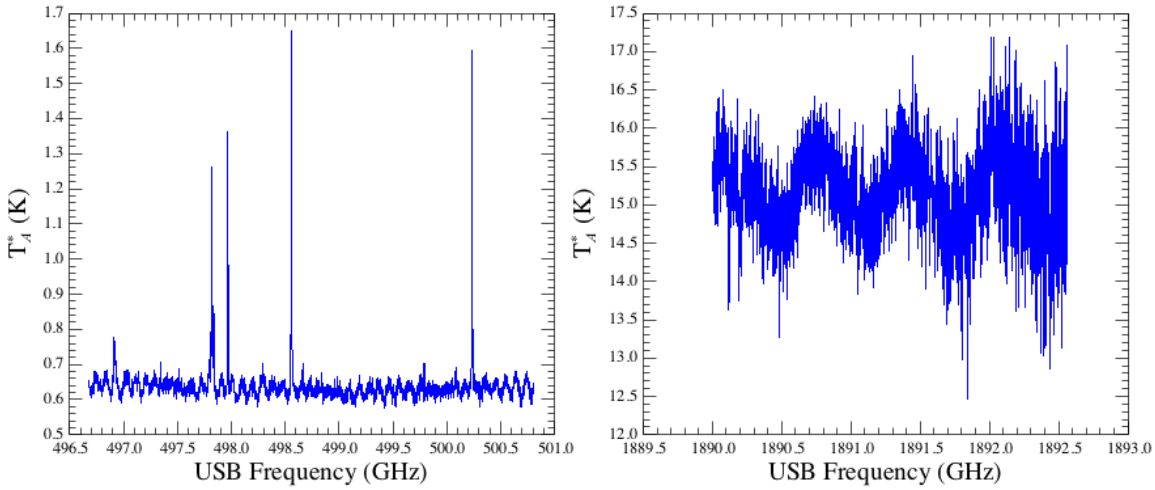

Figure 5.10: Example of strong optical standing waves. Left: optical standing waves from the internal loads in band 1a, with period near 100 MHz. Right: optical standing waves from the diplexer roof top mirror in band 7b, with period near 645 MHz.

5.3.1 Baseline distortion and residual standing waves

Optical standing waves

Optical Standing Waves (OSW) are resonances created in the cavity between two optical components of the HIFI focal plane or of the *Herschel* telescope. Although the optics design of *Herschel* and HIFI has given particular attention to the mitigation of such resonances, residual standing waves are still observed in some of the observations in the form of sine modulations. While their period is well defined by the size of the cavity where they originate, their amplitude will depend on several parameters such as instrument gain drift magnitude or the intensity of the signals entering the resonance.

The main optical standing wave components affecting the HIFI data have been compiled in Rodríguez-Coira & Teyssier (2015). They depend on the band in use but the following general rules apply:

- the most common optical standing wave in the HIFI data has a period near 100 MHz. In the low frequency bands it originates mostly from the internal load (it will in fact appear as a beat between the standing waves of the two loads of periods near 92 and 98 MHz respectively), while in the diplexer bands such as band 3 and 4, this standing wave originates from the LO feed horns and is most prominent at the IF band edges. The left panel of Figure 5.10 illustrates how such a standing wave can manifest in the top level products.
- HEB bands are mostly affected by a so-called Electrical Standing Wave (see thereafter), but they will also suffer from a 645 MHz period standing wave originating in the roof-top mirrors of the diplexer optics (Section 2.3.2). This latter also exists in the diplexer bands 3 and 4. The left panel of Figure 5.10 illustrates how this particular standing wave can manifest in the top level products.

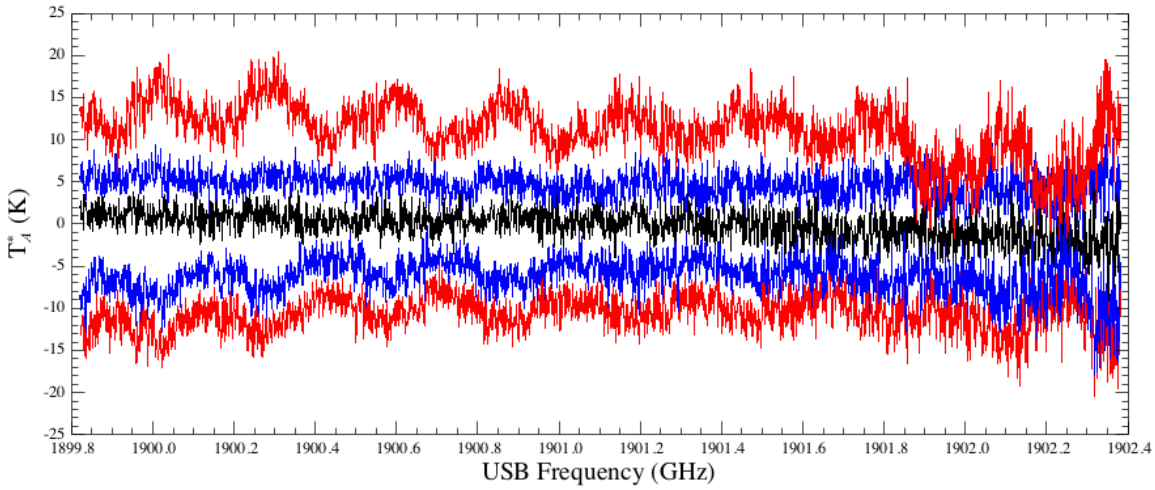

Figure 5.11: Example of Electrical Standing Wave in a band 7b observation. The various spectra correspond to different positions on an OTF map. Note the correlation between standing wave amplitude and continuum offset, as well as between the standing wave phase and the sign of the continuum offset (see Figure 5.12).

- the amplitude of the optical standing waves coming from the internal loads is multiplied by that of the incoming source signal (see Eq. 4.1) so that sources with strong continuum will feature very large standing wave amplitudes. Fig. 6.3 illustrates one example of such extremely enhanced baseline ripple. Mitigation measures against this effect are discussed in Section 6.4.1.
- **the fact that standing wave modulations are not featured in a spectrum baseline does not imply that no standing wave affect the spectral lines, rather this is indicative that the observed source has a negligible continuum. Optical standing waves are indeed multiplicative in nature and therefore apply similarly to any spectral channel featuring noticeable intensity.** The typical magnitude of the corresponding residual modulation is given in Section 5.8.2.

The main approach to correct the OSW consists in fitting sine wave models to the baseline. Several packages exist in various software distribution. Section 6.4.2 will discuss the mitigation of these artefacts in the framework of HIPE using the *FitHifiFringe* task (see also the HIFI Data Reduction Guide).

Electrical Standing Waves in HEB bands

Electrical Standing Waves (ESW) are caused by reflections inside the coaxial cable between the HEB mixer and the first signal amplifiers (Higgins 2011 and Higgins & Kooi 2009). As such the ESW are created downstream of the mixing process and their shape does not depend on the LO tuning frequency. They depart from OSW in that they cannot be defined analytically by a single amplitude or period. In the Fourier space they show up at a period of about 320 MHz, which is a good approximation for first order correction with an OSW model over narrow frequency ranges.

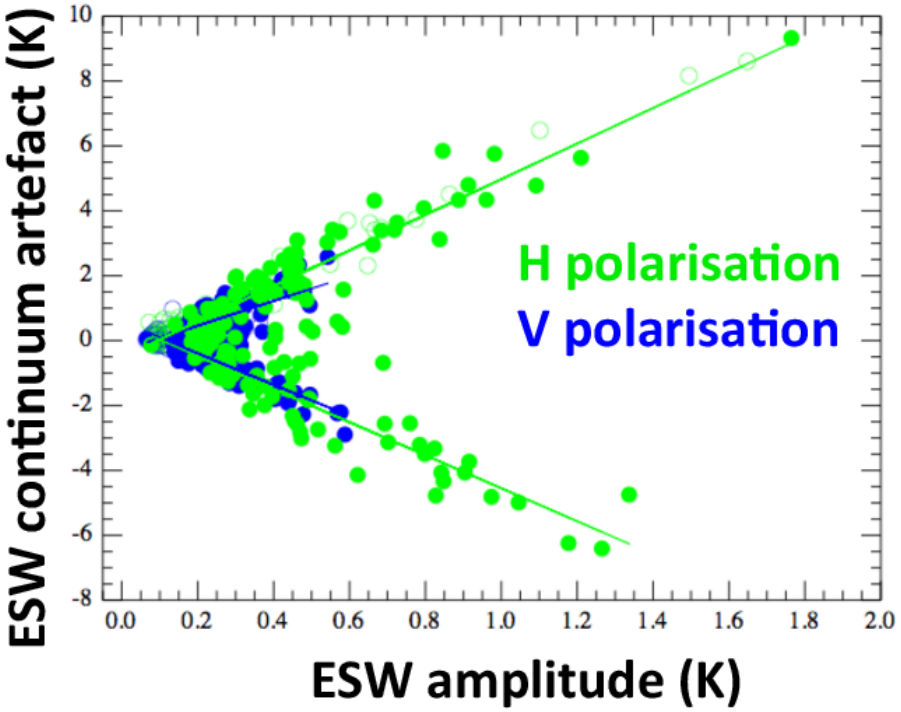

Figure 5.12: Correlation between ESW continuum artefact and ESW amplitude (H: blue, V: green) – the latter is here measured as the 1-σ noise RMS of the spectrum. The slope of the relationship between the two quantities is typically $\pm$ $\sim$5 depending on the modulation phase.

Fig. 5.11 shows some examples of ESW. Another peculiarity of ESW is that they contain a continuum component that scales together with the amplitude and phase of the modulation. When the phase turns by 180 degrees, the continuum component is actually negative. This means that the continuum information of observations taken in HEB bands cannot be fully trusted as long as some ESW is still present in the data. Fig. 5.12 illustrates the relationship between ESW amplitude and continuum contribution. The measured slope can be used to derive the typically residual continuum artefact based on the observed residual ESW amplitude in a given spectrum.

A dedicated algorithm was developed in order to correct the ESW (Kester et al. 2014) and is applied within the pipeline - see Section 12.4 of the Data Reduction Guide for more details. One should however be warned that residual ESW effects may still be present in a minority of observations, most likely in On-the-Fly maps taken in bands 6 and 7 (Section 5.5.3).

Baseline distortion at IF band edge

The Allan variance times have been observed to be usually poorer at the edges of diplexer bands. This leads to potentially degraded baseline quality as the drift may not be optimally compensated for in those areas of the Intermediate Frequency. This effect can worsen in case the LO injects significant excess noise in those areas. Such effect is illustrated in Figure 5.13.

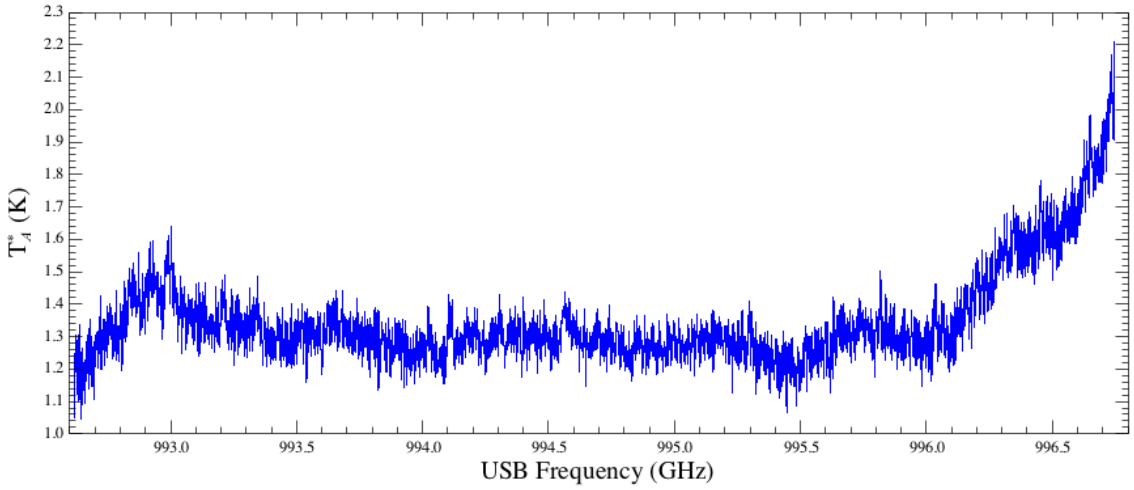

Figure 5.13: Illustration of strong baseline distortion at the edge of the IF range for a diplexer band (4a).

There is no real ad hoc mitigation measure other than attempting to flatten the baseline using polynomial or spline baseline fit models (see Section 6.4.2).

Platforming

It has been observed that observations with the WBS can suffer from platforming (also known as *stair-casing*) between neighbouring subbands of this spectrometer. This manifests itself as apparent continuum level offsets between adjacent subbands, that can be significantly larger than the noise levels. About 10% of all HIFI observations have been identified as being affected by this artefact, and the vast majority of them have been taken in the Position Switching mode. Fig. 5.14 shows some examples of such effect. A dedicated quality flag has been introduced in order to warn the user against possible platforming (see Section 10.4 of the HIFI Data Reduction Guide). It reads ”Platforming present in overlapping subbands” and will be present in the quality summary when relevant.

Correcting for platforming basically means re-adjusting the various subband levels to a common level. Assessing the ”true” level is not obvious, however, some general rules can be followed:

- it is usually the case that one particular WBS subbands deviates from the overall level featured by the other subbands. As such it is a fair assumption that the baseline level of the majority of the subbands can be used as the reference level
- the HRS is not affected by platforming, so in case of overlap with the affected WBS subbands, it can be used to assess the reference level

Note also that subbands are no longer distinguished in top level products because those are stitched (Level 2.5 – see Section 6.1). Strong platforming will appear as significant discontinuities over the spectrum. Correcting from such an effect is therefore best done on non-stitched data, meaning Level 2 products (Section 6.1).

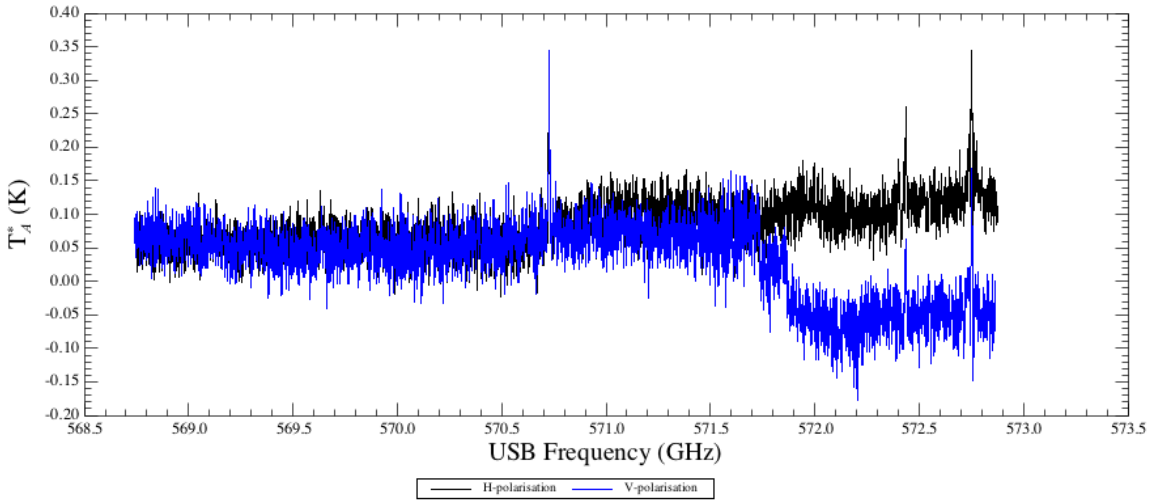

Figure 5.14: Example of platforming in the upper end of the V polarisation spectrum (the H polarisation spectrum is here not particularly affected by this effect.

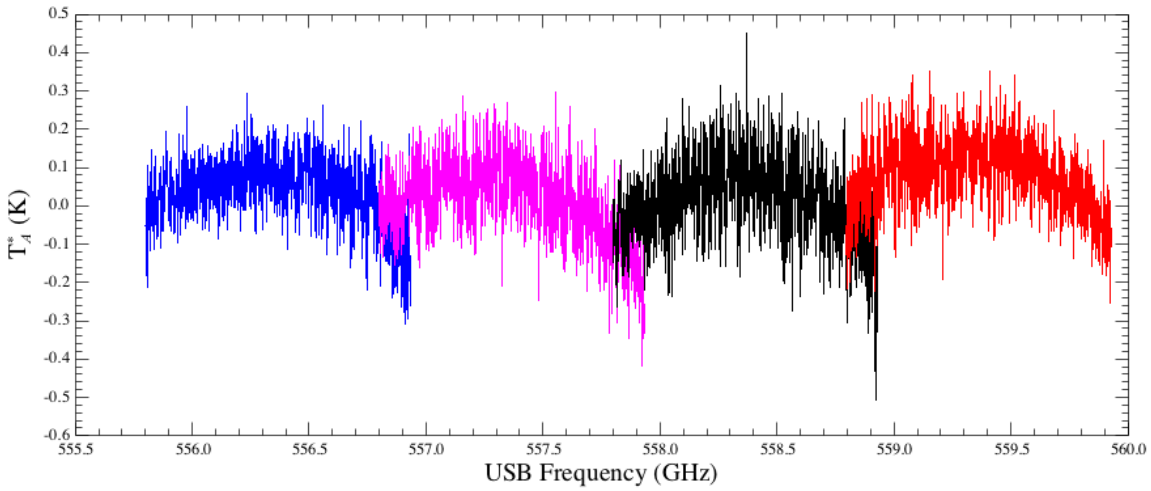

Figure 5.15: Example of anomalous parabolic baseline shapes in Wide Band Spectrometer data.

Parabolic baseline shape in PSW and OTF-PSW

Observations taken in Position Switching (PSW) and On-the-Fly (OTF)-PSW modes can sometimes suffer from residual baseline structures of parabolic shapes. Those are seen to only affect WBS data, and are illustrated in Fig. 5.15. The best way to mitigate those is to fit a second-order polynomial to the individual subbands - like for the platforming (previous section), this is best done on non-stitched data, and comparison to the HRS data may provide additional information concerning a reference level for the uncorrupted baseline continuum.

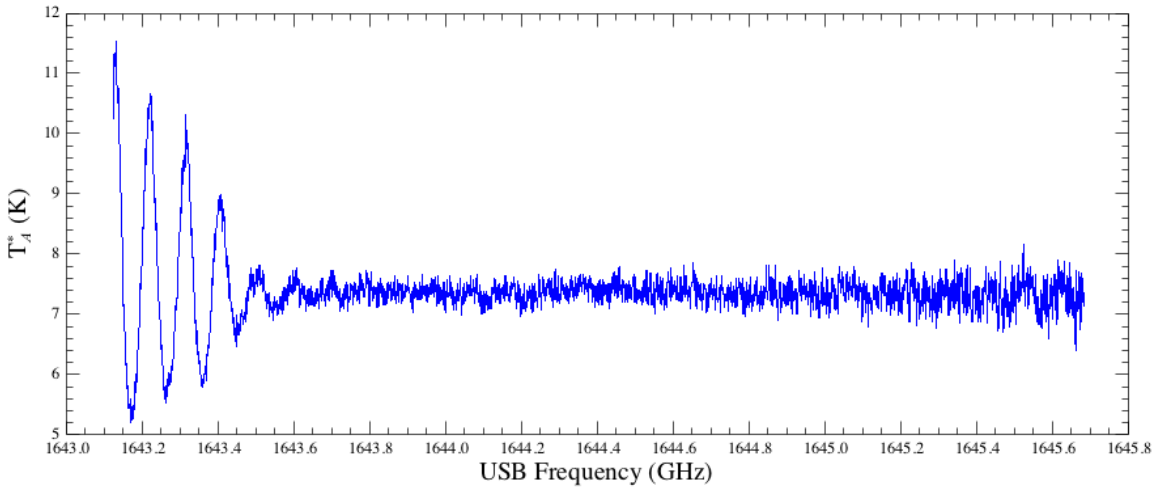

Figure 5.16: Example of strong oscillation in the lower end of the IF in band 6a.

Table 5.2: Band 6 resonances.

Fast Oscillations in band 6

Data taken in bands 6a and 6b are affected by significant baseline oscillation in certain ranges of LO frequency tunings. This ripple is generally limited to the lower part of the IF (typically the first 300 MHz) and has an apparent period of 100 MHz. At this stage the origin of this artefact is not clear but it is seen predominantly in the load data, and gets propagated to the science spectra during the bandpass calibration. Fig. 5.16 illustrate spectra affected by this modulation. Table 5.2 summarises the LO tuning ranges where this effect occurs.

There is no easy correction measure known to date, so that it is recommended to discard this frequency range. It has, however, been noted that changing the smoothing width of reference measurements in the case of Load Chop observations has improved the situation (Goldsmith et al. 2015). Please refer to Section 6.4.1 for a description of how this could be done.

Slow LO settling times

There is an intrinsic thermal settling time involved in the HIFI electronics every time the instrument has to be configured to observe at a new frequency. In Spectral Scans, these re-tunings occur at a relatively quick pace since only limited time can be afforded at each frequency point (typically up to 10 seconds ON-target per LO tuning). Depending on the frequency, it may happen that the settling time exceeds the dead time allocated to the tuning block before starting the observation. In this case, part of the spectra may be acquired under sub-optimal stability conditions and result in imperfect baseline calibration at the Level 2. This rare problem usually occurs at isolated ranges of a given band, as illustrated in Fig. 5.17. Note that such baseline jumps should not be present in Spectral Scans for which *Highly-Processed Data Products* (HPDP, see Section 6.2) will have been generated.

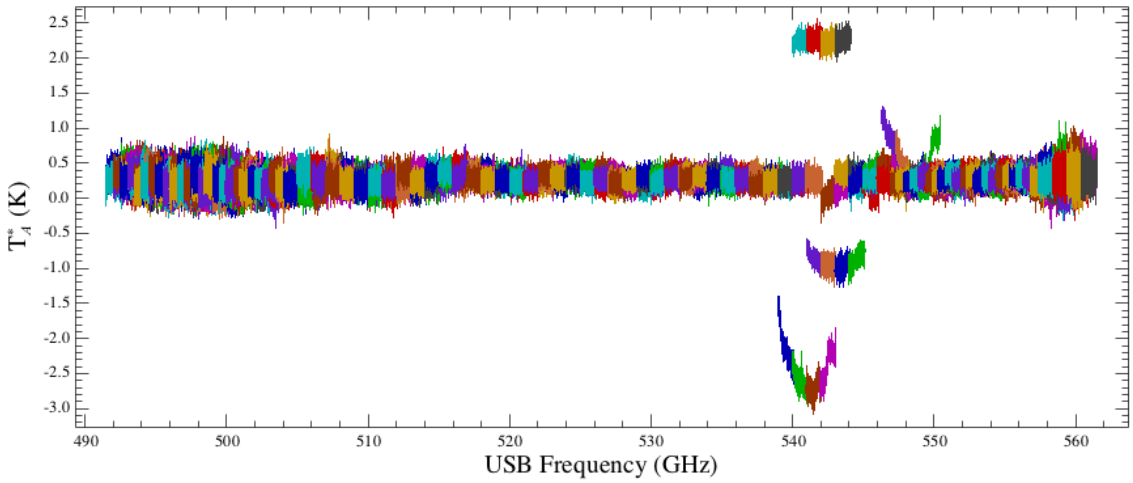

Figure 5.17: Illustration of baseline jumps due to slow LO thermal settling at isolated tuning frequencies.

One possible mitigation consists in throwing away part of the Level 2 data that were taken when the system was still thermalising. This is only possible if Level 2 data involve the average of several datasets. For this, the HSA will also provide non-averaged Level 2 data for all point-mode observations as *Ancillary Data Products* (see Section 6.2).

5.3.2 Spectral purity issues

The spectral performance of HIFI is intrinsically related to the quality and stability of the Local Oscillator signal. In order to allow heterodyne detection in one single pair of Lower and Upper Sidebands (LSB and USB), a source signal is amplified and multiplied to the desired sky frequency in one predominant frequency tone, providing therefore a mono-chromatic reference signal. It can happen that the multipliers, which are basically diodes, enter into oscillation and make the LO signal multi-tone. In this situation the heterodyne detection also takes place in other spurious sidebands.

It has been observed since the early pre-launch tests of HIFI that its LO chains would be locally affected by such oscillations, leading to a degraded spectral purity. Details about the test campaigns and correction measures taken over the various phases of the pre-launch and orbital life of the instrument are given in Teyssier (2017c). We summarise here the impact on the archival data and give recommendation on how to deal with the affected data.

Spectral impurity will manifest in the data in two main forms:

- *Spurious signals*: these correspond to spectral features created inside the instrument. The most common forms is what we call *spurs*, that mostly show up as more or less narrow spectral glitches in the data. Because those glitches slowly move in frequency with time, in the Level 2

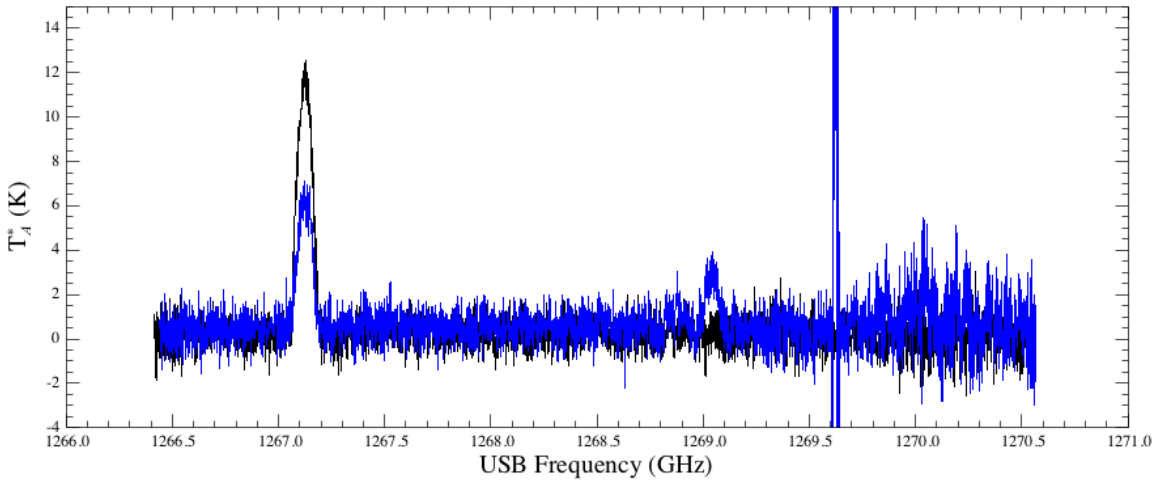

Figure 5.18: Example of spectral spurious response in band 5b. The black spectrum represents data obtained with a spectrally pure LO. In the blue spectrum, affected by spectral impurity, the intensity of the line at 1267.2 GHz has an erroneous intensity calibration, and a prominent ghost line exists at 1269 GHz, on top of narrow spurs at 1269.5 GHz and LO excess noise at the spectrum upper end. Note that such impure data have not been publicly released in the HSA.

data they will appear as a combination of positive and negative components due to the subtractions/divisions involved in the pipeline. There exist broader spurs, that look like bumpy baseline structures, and the most drastic version of them will cover one or more spectrometer subbands and manifest as saturation due to their very large magnitude. In the case of strong narrow spurs, cross-talk between the WBS CCDs can lead to the spur being mirrored regularly in each of the spectrometer subbands. Figure 5.19 illustrates some of these cases.

- *Spurious response*: this corresponds to LO tunings where the instrument is sensitive to more spectral ranges than the solely intended Lower and Upper Sidebands. This will manifest both as ghost lines coming from other frequency ranges, as well as an erroneous line intensity calibration since the measurement on the internal load will also hold broadband signal arising from more than the sum of the respective LSB and USB. Fig. 5.18 illustrates this situation in band 5b.

It is possible to ”purify” some of the affected LO tuning ranges by making specific adjustment to the LO chain multiplier settings. This was achieved in several HIFI bands, as described in Teyssier (2017c). This means that purity issues in some frequency domains are time-dependent. We emphasise in the following sub-sections some prominent spectral purity issues that applied to the flight data.

Spurious response

Table 5.3 provides a summary of the most relevant ranges affected by spurious response. The worst case was that of band 5b, which was so beset by spectral impurity that it had to be withheld for scientific usage until OD-595. We should also warn against some isolated areas of band 5a (mostly

LO band	LO Frequency range (GHz)	Applicable period (OD)	Affect on calibration
2a	713–715	Whole mission	Line ghosts
3b	951.3–952.9	until OD-1396	Unreliable line intensities
	952.9–953	Whole mission	Unreliable line intensities
5a	1231.7–1236	until OD-595	Unreliable line intensities
5b	1236–1272	until OD-595	Line ghosts and unreliable line intensities
7a	1713–1718	Whole mission	Negligible
	1756–1758	Whole mission	Negligible
7b	1866–1888	until OD-305	Negligible

Table 5.3: Summary of most relevant spectral ranges affected at one point of the mission by spectral purity issues.

at its upper end), which were released for use early in the mission, but had to be purified later on. In those ranges, the line intensity calibration can be erroneous prior to the purification campaign. Finally it should be noted that spurious responses have been detected in bands 7a and 7b, however, they did not seem to have detectable impact on the intensity calibration. They are still listed here for the sake of completeness.

Spurious signals (spurs)

The list of all recorded spurs cannot be given in an exhaustive fashion in this document. Instead, the HIFI ICC has attempted to compile all detected spurs in dedicated tables stored in the Calibration Tree (see Section 6.5). For each observation, a list of spurs is provided, and used by the pipeline in order to appropriately flag the data (use of flag `SPUR_CANDIDATE` or `IGNORE_DATA`). In Spectral Scans, these spur tables have been built manually by Calibration Scientists and allow the generation of spur-free deconvolved spectra. When spurs are excessively strong and/or broad, they will appear as saturation, as illustrated in the lower right panel of Figure 5.19

There is however a handful of spurs that got particular attention during the mission, typically due to their proximity to some of the main science target lines of HIFI. For the majority, correction measures were taken during the mission in order to clean the spurious signals from the affected bands. We list thereafter some of the most noticeable features present in the archival data.

Band 1a spur:

A very strong spur did affect the band 1a data in the LO frequency range [535–553] GHz. In some tunings, the spur was so strong that it lead to partial or full saturation of some of the WBS subbands, but also was present as a weaker cross-talk in the other 3 WBS subbands. The upper left panel of Figure 5.19 illustrates some of these effects. This spur received particular attention early on because it did affect tunings targeting the ground-state of the *ortho*-H_2O line around 557 GHz. The problem lead to most users re-designing their AORs in order to pick up the line in the LSB of band 1b (together with NH_3), however a handful of observations did stick to the band 1a tuning. In those latter cases, narrow water lines can be salvaged from the largely destroyed baseline but special care needs to be taken there.

The affected frequency range was cured and from OD-496 onwards the spur was no longer present in the data, with the exception of tunings at LO Frequencies around 542.3 GHz and above 551 GHz.

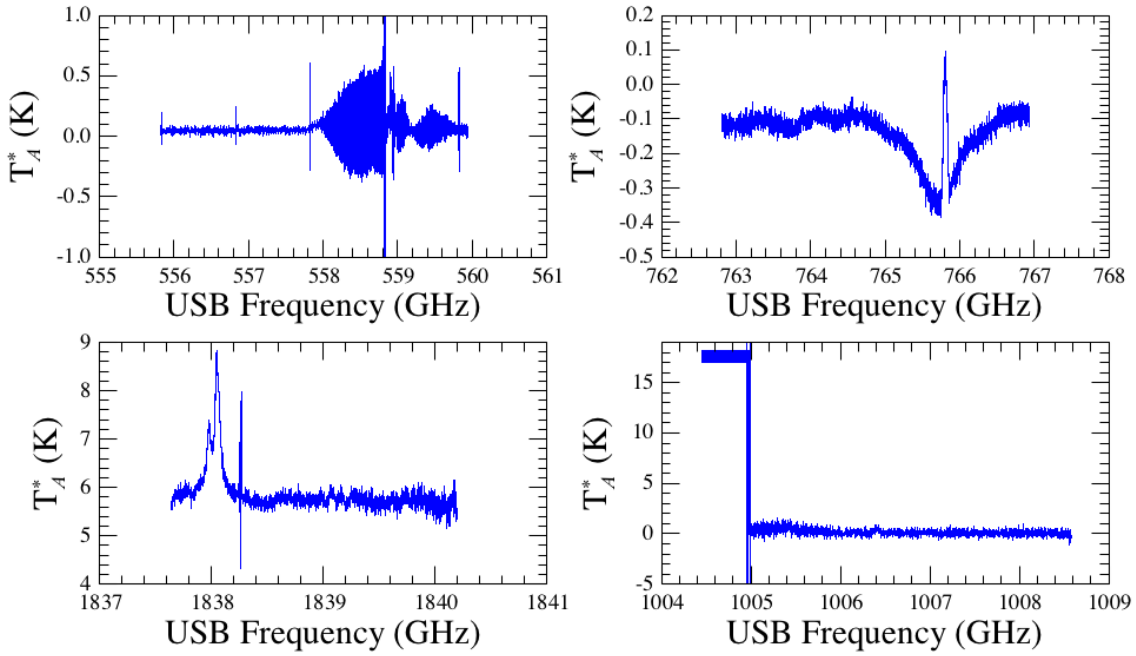

Figure 5.19: Example of strong narrow and broad spurious signals. Upper left: spur present in band 1a for data taken up to OD496. Upper right: example of broad spur present in band 2b. Lower left: spur present in band 7b for data taken up to OD-779. Lower right: example of saturation due to very strong spurious signals. See text for further detail.

Band 2b broad spur

A broad (about 1 GHz width) spur has been observed in LO frequency tunings around 752 GHz. It became particular relevant because this setting was frequently used to observe the $2_{11} - 2_{02}$ transition of *o*-H_2O. The upper right panel of Figure 5.19 illustrates how the spur is intermittently present depending on the precise LO tuning – note the possible ambiguity it creates on the water line profile interpretation, especially in sources with strong continuum susceptible to show up in absorption. Because this spur is so large, and it overlaps with a valuable water line, it is not flagged in the default spur table. Instead it is recommended to the user to fit a typically second order polynomial to get rid of the distorted baseline.

Band 7b spur

A spur was regularly observed at LO frequency tunings in the range 1829–1839 GHz, which was particularly relevant for observations of the OH line near 1838 GHz. The lower left panel of Figure 5.19 illustrates this spurious signal. The spur was cleaned from OD-779 onwards.

5.4 Anomalous data and observational events

On very rare occasion, some sporadic malfunction of the On-Board Software have lead to data corruption or failed commands. These resulted in anomalous data at Level 0 that propagated up to the

Level 2 and Level 2.5 products. Correction measures have been applied to the pipeline, so that such effects should be invisible in the archival data. However, for expert users willing to work from the low level data products, some precautions need to be taken. Some more details about those events and recommendations on how to handle them are given in Appendix B.

On top of that, some HIFI observations were affected by Single Event Upsets (SEU), which are in effect high energy particles hits onto their on-board control units. Those lead typically to the interruption of the observation (the first recorded HIFI SEU on OD-81 actually killed the prime unit of the HIFI instrument, see Jellema et al. 2010), and making the corresponding observation failed. Since failed observations could still provide useful data to expert users, we also provide in Appendix C more background about those events.

5.5 Validation of the Observing Modes

This section describes how the HIFI Observing Modes have been validated, and their performances characterised with an emphasis on the aspects that are relevant to the analysis of observations in the HSA. An understanding of each mode's performances has relied upon a complete examination of instrument and telescope interactions under a variety of set-ups contained in the AORs on diverse celestial targets, and upon the state of the data processing. The principle metric for validating an Observing Mode is the baseline noise which is achieved compared to the modelled noise performances. Portions of this chapter incorporate information which was written into the HIFI AOT Release Notes (Morris 2011), as well as HIFI Information Notes on the the HIFI Calibration Web page concerning key calibration updates and specific effects in the Observing Modes that the user should be aware of. The basic features of the Observing Modes available with the three HIFI AOTs are summarised in Chapter 3.

5.5.1 Validation Approach

The strategy for validating the HIFI Observing Modes was planned to be carried out over the Performance Verification (PV) phase in a manner typical of most astronomical instruments – establish specific functional and scientific performance metrics and observe a variety of celestial targets with each of the modes in the parts of the sky that are available to the telescope over the validation period, including but not limited to calibration sources both compact and extended, and for which similar observations already exist (e.g., ground-based) or would be carried out in the Key Programmes. The expected performances of the Observing Modes would be validated with different AOR setups on different celestial objects, but also on the same object and target lines for Mode cross-comparison and instrument stability tests. The availability of comparable ancillary observations was not a strict requirement since mode performance can be checked based on functional and noise-related metrics, and in fact *Herschel's* launch in May 2009 resulted in PV being scheduled in a visibility window with very limited availability of the most well-known sub-mm sources.

The AOT validation approach was drastically revised, however, as a consequence of the Single Event Upset (SEU) that occurred on August 3, 2009 (Jellema et al. 2010), resulting in the loss of the Prime Side electronics sub-system, and delaying the resumption of scientific operations until late January 2010 on the Redundant Side electronics. The nominal AOT commissioning strategy was revised from a scheme of Observing Mode validation in which the Modes would be validated essentially in parallel,

to one which serially emphasised the most requested modes in the Key Programmes, under the notion of a potentially limited instrument lifetime and sudden termination of operations. This accelerated the release of the modes according to their subscription, basically ordered as follows:

1. Dual Beam Switch modes at normal and fast chop rates, all three AOTs.
2. Position Switch, Frequency Switch, and Load Chop modes for Single Point and Spectral Scan AOTs.
3. Position Switch, Frequency Switch, and Load Chop modes for OTF Mapping.

Users of the HSA will find a frequent number of observations which were carried out for the dual purpose of AOT performance testing and providing spectroscopic data for early "priority science" which had been formulated between the HIFI Project Scientist, Key Programme PIs, and the AOT/Uplink team in the period leading up to resumption of operations. The release of each Observing Mode was thus balanced between scientific urgencies of the Priority Science Programme, and iterative completion of fundamental operating procedure optimisations for, e.g., the mixers and LOs in order to provide the purest stable signal at the lowest system temperatures possible. Fortunately, celestial source availability became more favourable with access to targets such as Orion, the Galactic Centre and other heavily-subscribed KP targets when AOT testing resumed in February 2010, providing a faster path to HIFI's scientific results. After the initial release of all Observing Modes requested in the KPs, the HIFI-ICC continued to characterise performances into the Routine Phase, incorporating experience with instrument purity and stability, updates to the calibrations of the mixer beams, characterisation of the standing waves and their effects on observations by mode, and maturing data processing software. For the remainder of this chapter, we attempt to concisely summarise these evaluations as most relevant to users of the HSA.

5.5.2 Mode Usage Guidelines During the Mission

At the time that each Observing Mode was validated and released for planning and scheduling, certain caveats and recommendations were communicated to users through the HSC. Summarising those aspects which remained in place for portions or all of the mission, and which affect the complexion of observations in the HSA:

1. DBS Modes for All AOTs.
    - In HEB bands 6 and 7, users will find most DBS Point, DBS Spectral Scan, and DBS Raster observations in the HSA to have been carried out with the fast chop option. This was recommended by the ICC, following comparisons of data using the normal (or slow) chop DBS with FastDBS modes, indicating that the latter generally perform better with respect to correction for electrical standing waves and baseline drift (Section 5.3.1).
    - For mapping observations earlier than obsid 1342200757, AORs which requested Nyquist sampling use a definition in which the spacing between map points is actually "half-beam" spacing, i.e. HPBW/2.0. After this point, the Nyquist definition was changed to HPBW/2.4, and the Half Beam spacing option was subsequently offered in HSpot. Only the earlier OTF maps and DBS/FastDBS raster maps with the `nyquistSel` keyword set to "true" in any of the spectral data cube headers are affected. This change was documented at the time in the following technical note.

- DBS/FastDBS-Cross mapping modes, which offered two single scan lines orthogonal to each other, were not generally used over most of the mission. The mode was removed from HSpot v6.0 and later versions, due to a timing error which could not be isolated in the instrument, telescope, or interactions between the two, resulting in data in the ON and OFF phases being disordered.

2. Frequency Switch and Load Chop modes with all AOTs.

   - The Frequency Switch mode was initially available in all LO bands for all three AOTs, but was discouraged in Bands 6 and 7 due to indications of poor stability. The mode was discontinued for Spectral Scans in Bands 6 and 7 early in the Routine Phase. The difficulty was to allow sufficient time for settling of LOs at the frequency throw under the constraint of generally very short Allan times, and the occasional unstable pairings of the reference measurements after retuning while mixer currents are drifting, a particular problem for Spectral Scans. The wait times needed to mitigate these effects were judged to have unacceptable impact on observing efficiencies. The lack of stabilisation can be manifest in the data as baseline drift and non-linear distortions, strong standing waves, and artefacts that may not be treatable in interactive data processing, such that achieving the SNR goals set by the user in HSpot could not be guaranteed within acceptable margins. Only six Spectral Scans using the standard (non-engineering) Frequency Switch mode in Bands 6 and 7 for science are in the HSA, and most are released to the public. However some of these are marked as failed in the Quality Control Status column of the query output table, thus only the products manually-processed by the HIFI ICC as so-called "Highly-Processed Data Products" (HPDPs – see 6.2.3) are recommended. Similarly, only HPDPs for OTF maps using Frequency Switching in Bands 6 and 7 are recommended.
   - The Frequency Switch and Load Chop modes are normally used with a position switch to a line-free sky position selected by the observer, normally intended to correct the baselines for standing waves. However this option could be skipped for sources where the observer expected very strong spectral features and was not concerned with the standing waves in the baseline, or a suitable line-free position for the reference measurements could not be found. These kinds of observations are identified in HSA query output with "`NoRef`" attached to the end of the Observing Mode name. If the NoRef observations are have not been improved as an HPDPs, then the spectral output of the SPG should be expected to contain optical standing waves at their maximum amplitudes in all bands. The electrical standing waves in Bands 6 and 7, though, will have been treated in the pipeline.

3. Special Notes Applying to OTF maps.

   - OTF mapping with Frequency Switch was discontinued in HSpot in Bands 6 and 7 for reasons of poor stability as described above. The exception was over the range of LO frequencies 1890 - 1898 GHz which was optimised for stability and purity in the area of the C^+ 1900.537 GHz line, for low-redshift sources. All scientifically serviceable OTF Frequency Switch maps are provided in the HSA as HPDPs.
   - OTF maps using Frequency Switching or Load Chop do not have even sampling in each scan line, accounting for the periodic baseline calibration measurements (either a measurement at the frequency throw or internal load measurement) taken while the telescope is moving. The sampling requested by the observer is met *on average*, meaning that the

ON-source map points are more closely spaced than requested to compensate for the reference "holes". In the HIFI pipeline, the spectral cubes created from OTF maps with these modes use an extrapolation (or interpolation) of map point signal that avoids sampling noise (NaNs) at the positions where internal measurements were taken.

- The line-scan-with-off telescope pointing mode used to execute all OTF maps with position switching was found to deviate from the expected path of the line scanning in a way that produced significant non-Gaussian pointing noise in the data (inaccurately but historically referred to as OTF "zig-zag"). Final improvements in the reconstruction of the telescope's high-frequency pointing drift while observing does help to mitigate the effects on the data, with more accurate WCS-based weighting of map point signal in the convolved cubes. However the effects are not entirely removed in individual (non-coadded) observations, generally, since at each map point the effects of large path errors by more than 15% of the beam size on the a high-signal target cannot be corrected for an arbitrary astronomical source. The effects, however, are fully untreated only for a small number of OTF observations, for which improved pointing reconstruction could not be applied.

5.5.3 Noise Performances

In this section the noise performances of the Observing Modes are summarised, highlighting effects which the user may find helpful when examining HIFI products in the HSA. Spectral line and baseline noise properties are key measures of how well the AOT parametric assumptions contained in configuration and calibration tables capture the performance of the actual instrument and telescope systems. The results must reflect acceptable margins of error on the approximations that are necessary when time estimation in HSpot was carried out by the observer.

Predicting and measuring baseline noise in HIFI spectra

Figure 3.3 gives a schematic representation of the interfaces involved in the estimation of noise and timing parameters for an AOR created by the Observer, and for which upper and lower spectral resolutions have been entered to set the desired noise level range that meets the science goals of the planned observation. The role of the desired minimum and maximum spectral resolutions set in HSpot is important: all observations are taken at the *native* resolution of the instrument, which for the WBS is $\simeq$1.1 MHz, but depending on the sensitivities at the LO frequencies of interest, the baseline noise at the highest resolution possible for the backend spectrometer may never be reached in reasonable observing times. Thus as discussed in Section 3.2 the noise values provided in the output of time estimation in HSpot have been calculated at the two resolution bandwidths. The timing loops (illustrated in Figure 3.1) are driven by the minimum bandwidth (maximum goal resolution) entered by the Observer.

Note that the pipeline provides spectral products at the native resolution, with no smoothing applied. This is up to the user, which may have somewhat different goals compared to the original Observer. Noise is measured in the pipeline starting with Level 2 products (Section 6.1), using a special task (`mkRms`) that processes the data to intermediate levels in a manner consistent with the HSpot predictions including the smoothing.

To summarise the processing required to compare the noise properties of spectral products to the expected performances consistently:

1. In HSpot, the baseline noise over each of the two resolution bandwidths is predicted on a single sideband (SSB) $T_{\rm mb}$ scale, and assumes that spectra from both H and V mixers have been averaged. System temperatures measured and tabulated from separate calibration observations over each LO band determine the sensitivities.

   - In Spectral Maps and Point AOT observations, the noise at each bandpass resolution has been estimated either across the entire available IF, or within a 1 GHz window in cases where the Observer has chosen in HSpot to rely on an optimum WBS sub-band for both noise estimation and placement of the spectral feature of interest (the "One GHz option").
   - In Spectral Maps, the estimated noise applies to the spectrum at each map point, without accounting for the convolution of signal from neighbouring map points in the spectral cubes produced in the pipeline — thus the baseline noise measured at each map point should generally be lower than predicted.
   - For Spectral Scans, the noise is estimated at a reference frequency chosen in the AOT logic.

   The output of the time estimation, including the noise predictions and various observing time parameters, are stored with the AOR and propagated into the data products.[1]

2. To measure the baseline noise in the pipeline or in HIPE, consistent with the HSpot predictions, the Level 2 spectrum datasets from the H and V chains for all Observing Modes are first averaged together.

   - When more than one integration cycle has been carried out for Point modes, they have been averaged in the pipeline at Level 2, providing two H and V spectra to average together for noise measurement.
   - Similarly for Spectral Scans, if multiple integrations have been taken at each LO setting, the Level 2 datasets are averaged together, however then noise is measured from the H and V average of the two deconvolved (SSB) Level 2.5 spectra.
   - In Spectral Maps when more than one map cycle has been carried out, no averaging of the Level 2 datasets from the individual H and V dataset is done. At Level 2.5, the `doGridding` task grids the data for each polarisation into cubes, weighting the signal according to the astrometry of the input datasets, and the H and V cubes can be averaged for noise measurement at each map point.
   - The Point and Map data taken at a single LO frequency are not deconvolved, so a DSB to SSB correction factor 0.5 (assuming equal sideband gains) is applied. The HSpot "One GHz option" is taken into account, so that noise is measured over a 1 GHz window of the optimum WBS sub-band defined for the LO band in use.

   The noise is measured from spectral H+V averages, smoothed to the resolutions given by the minimum and maximum resolutions requested in HSpot, and also at native resolution. The resulting measurements from `mkRms` in the pipeline are stored as metadata in the observation context and underlying products (see Section 6.1).

[1] Due to limitations on the architecture of the HIFI uplink system in the first half of the mission, the HSpot noise and observing performance parameters may be incomplete in products with obsids earlier than 1342227392.

Table 5.4: Key HSpot time estimation output parameters.

Description	Parameter	Unit
Predicted SSB Noise USB at minimum bandwidth	noiseMinUsb	K
Predicted SSB Noise USB at maximum bandwidth	noiseMaxUsb	K
Predicted SSB Noise LSB at minimum bandwidth	noiseMinLsb	K
Predicted SSB Noise LSB at maximum bandwidth	noiseMaxLsb	K
PredictedSSB Noise mimimum bandwidth	noiseMinWidth	MHz
Predicted SSB Noise maximum bandwidth	noiseMaxWidth	MHz
Temperature (main beam) at noise reference frequency	TmbReference	K
Noise reference frequency	noiseRefFrequency	GHz
Observing time	observingTime	s
On source time	onTime	s
Off source time	offTime	s
Overhead	overheadTime	s
Total time efficiency[a]	totTimeEfficiency	%
Total noise efficiency[a]	totNoiseEfficiency	%
Drift noise contribution[b]	driftNoiseContrib	%

[a] The observing time and noise efficiencies are with respect to an "ideal" instrument, using the radiometric equation applicable to each observing mode.
[b] Drift noise contribution is with respect to the total noise including the radiometric component.

An important point about noise measurement with `mkRms` in the pipeline is that an attempt is made to detect and mask spectral features, using some initial baseline smoothing, fitting, and subtraction steps that depend on LO frequency (e.g. the smoothing widths). This does not work perfectly for all observations, particularly for broad features with peak intensities that less than 2-3σ of the RMS noise at native resolution, or complex features with multiple peaks and/or absorption lines. The noise tends to be overestimated for such cases.

Observations

Most of the observations analysed for their noise properties and repeatability were obtained specifically for the AOT checkout tests during PV, using the standard observing modes offered in HSpot. Many observations taken in AOT checkout are part of a block, for example to evaluate the efficiency and calibration accuracy tradeoffs in certain bands with FastDBS versus normal DBS mode with overlapping ranges of chop rates, or with and without continuum timing. Such blocks of tests yielded recommended optimum AOR setups or practices for different observing scenarios, summarised in Chapter 3. Many AOT performance observations were also carried out using a "procedural" version of the Observing Modes, which allows a wider range of uplink and AOR setup parameter values for testing. These were used also most commonly for pointing tests, focal plane geometry calibrations, and beam characterisation. These"Proc" modes are explained further in Section 3.3.4. Performance evaluations also exploit a number of HIFI science observations in the HSA.

The noise performances summarised below are based on largely interactive inspections of calibrated spectral products from selected observations carried out through PV and the early Routine Phase. These are the results which have been presented in the AOT Performance and Release Note (Morris

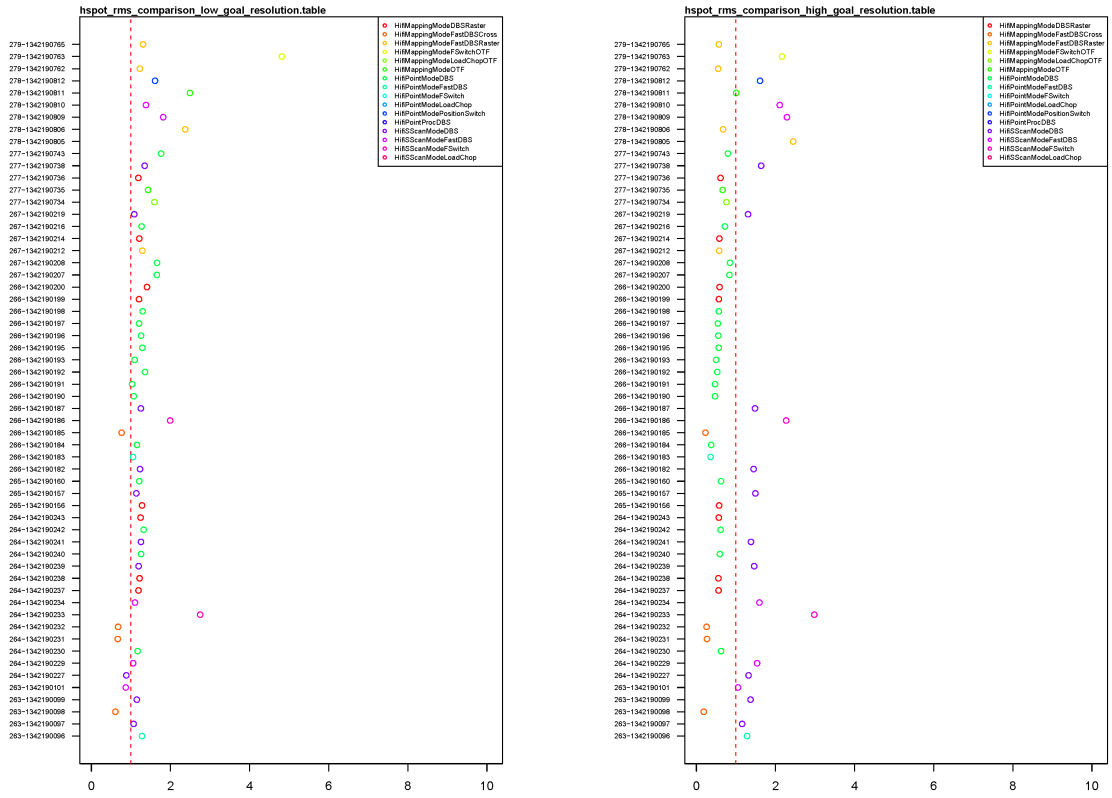

Figure 5.20: Ratios of measured to predicted (HSpot) noise for observations carried out through OD 263, with data smoothed to the goal low resolution (left) and high resolution (right). Y axis labels indicate OD-Obsid.

2011). The plotted points are from products produced in an early version of HIPE, using calibrations that do not exploit automated ESW corrections in Bands 6 and 7 (Section 5.3.1), or other advances in the pipeline. Since noise performance metadata are produced now for all HIFI observations in the HIFI pipeline (see Table 2), which also includes ESW corrections, we can in principle extend the evaluation of the latest products to cover the entire mission. As mentioned above, however, the automated line-masking needed to isolate spectral features from the surrounding baseline is not sufficiently robust against the richness of all HIFI observations in the archive. As a consequence we cannot rely on the raw auto-generated results, without following up individual outlier cases to arrive at reliable noise statistics. We give an example of this below. The full statistical view on AOT performances is presented in Morris & Ossenkopf (2017). The validation results summarised below are nonetheless representative of performance over the entire mission, based on spot-checking many observations, using different observing modes and with varied spectral complexity, over the normal course of calibration activities.

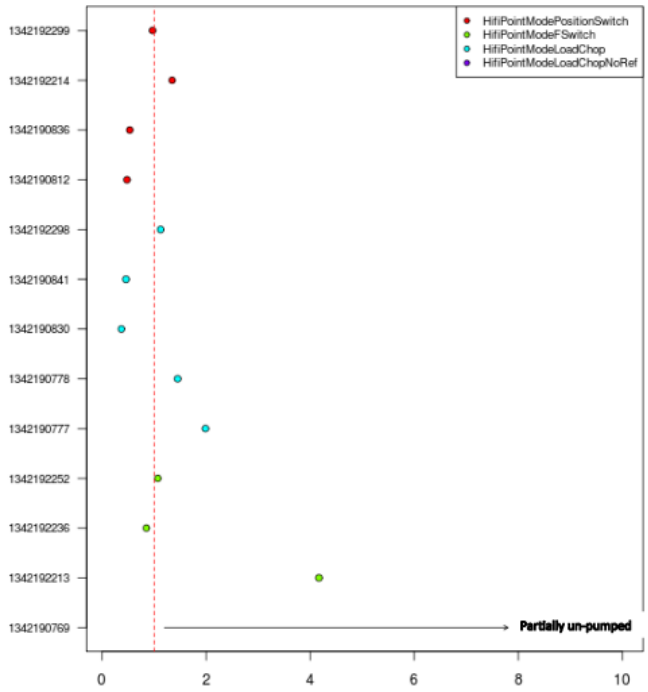

Figure 5.21: Ratios of measured to predicted (HSpot) noise for all Point AOT total power observations, with data smoothed to the requested high resolution. The largest outlier (1342190769) is a PointFrequency Switch observation in Band 7b with partially un-pumped data , prior to optimisations to increase instrument stability and tuneability over many frequencies in Bands 6 and 7.

Point Modes

Figures 5.20 and 5.21 show the ratios of observed over predicted noise, at the goal minimum and maximum spectral resolutions for each observation. In Figure 5.21 the outlier Point Frequency Switch observation (1342192213) represents a limitation to measure the baseline RMS noise when the spectrum is rich; see Figure 5.22. On the other hand, the outlier from observation 1342190777 was taken in Band 7b near 1897 GHz, where LO and mixer settings affecting purity and stability had not been finalised, and the data contained standing wave residuals of mismatched bandpass calibration during on-source and sky reference measurements. This became more of an issue for spectral maps and spectral scans due to their data collection timing constraints, as summarised thereafter. Aside from those data points, the measured baseline noise levels are generally within 30% of the predicted levels at both goal minimum (more smoothed) and maximum (less smoothed) spectral resolution.

Mapping Modes

Noise conventions

Referring to Figures 5.20 and 5.21, the DBS and FastDBS Raster mapping modes do not exhibit any remarkable differences with the corresponding Point DBS and FastDBS modes. However the

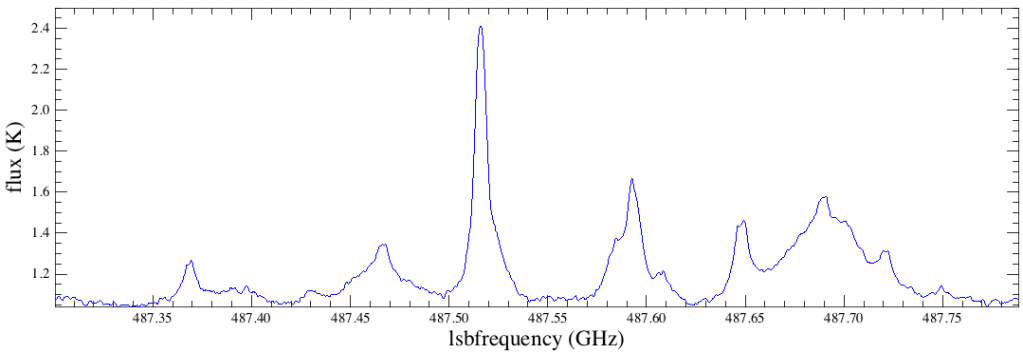

Figure 5.22: Obsid 1342192213, a Point Frequency Switch observation (unfolded) in Band 1a.

noise levels in both OTF and Raster maps are represented in a simplified way, in order to make the inter-mode comparisons, in that they represent the noise levels in all spectra in each mapped region averaged together. Thus they generally appear to show lower noise in the data than expected from the HSpot estimations. On an individual map point basis, the noise in spectra extracted from a processed spectral cube will be lower than the HSpot prediction, in general, since the predictions do not account for the convolution of signal from neighbouring map points. This implies that sampling will also have an effect on the measured noise: a Nyquist-sampled map will show a relative improvement of 18% compared to a Half-Beam sampled map of a region with a relatively flat brightness distribution. The chopped modes exhibit the most steady baseline noise levels across all mapping points, whereas the signal in individual readouts taken with the total power OTF modes can exhibit higher drift and accompanying standing waves (addressed below), especially in Bands 6 and 7. Edge effects may also be present in small maps regardless of Observing Mode.

In general, for comparison to HSpot predictions, the user should measure the noise in spectra extracted from individual map points of data gridded into spectral cubes where the number of pixels are precisely set according to the actual number of scan lines that were carried out. Alternatively one can use the spectra in the Level 2 products prior to gridding, remembering to combine the right number of datasets at each sky point according to the supersample factor and number of map cycles. Averaging these together will yield a crude means for comparing to the HSpot predictions. In either approach, the resulting RMS values should be on the optimistic side of the HSpot values if the form of the data (DSB, single polarisation) and measurement bandwidths are factored in consistent with the HSpot calculations. Based on the nominal results the user can then decide if regridding the data to higher samplings (smaller pixels) can be afforded.

Baseline drift and standing wave effects

In mapping observations, the timing of calibration measurements (internal and on the sky) for offsetting the effects of instrumental drift involves interleaving visits of the HIFI chopper mirror to an internal thermal load after a number of map points within each scan leg, and manoeuvring the telescope to a sky reference position between intervals of scan legs. The DBS and FastDBS raster modes utilise chopping rates which generally compensate for the drift, which can be seen by the usual lack of standing waves in the calibrated spectra. The noise performances of the DBS modes are similar in the Point and Mapping AOTs, as shown in Figure 5.20.

For the total power modes, the telescope may not be beam-switched to the sky reference position,

fast enough in bands where stability times are short, so that the final spectral products are susceptible to artefacts related to uncorrected (or not fully corrected) drift. This is mostly expected in the total power mode of standard OTF maps, where LO thermalisation after switch-on is more important than in the chopped modes (Section 5.3.1), the reference position may be up to 2 degrees away from the map centre, and the HEB bands 6 and 7 in particular have very short Allan times (Section 5.3.1 and Figure 5.11).

Band 6 and 7 pipeline products benefited from SPG version 13.1 from an automatic ESW correction that drastically improved baseline quality, and thus the net noise performances, for most observations in these bands. It is not uncommon, however, that a few datasets within an arbitrary observation could not be well-corrected for ESWs. Furthermore, residual optical standing waves left after the standard reference subtraction in the pipeline are still untreated in the end products. In cases still affected by electrical and optical standing waves and associated baseline drift, the gridded spectral cube will show detrimental effects as convolved instrument artefacts, leading to poor noise performances. The OTF Load Chop mode was therefore recommended to users in Bands 6 and 7, where test observations had shown similar extrema of drift levels and standing wave amplitudes in individual spectra, but overall a lower RMS of all spectrum baselines around a nominal level. On the other hand, exceptionally low stability of Frequency Switching in OTF maps compelled the ICC to forbid this mode in Bands 6 and 7, except over a range of most stabilised LO frequencies that allows placing the C^+ line at its rest frequency in the upper sideband (Section 5.1 and 5.5.2).

As mentioned in Section 5.3.1, the baseline offsets due to the ESW will scale with the ESW amplitude itself, as illustrated in Figure 5.12. In case of sub-optimal correction of the ESW affect, this has direct consequences on the morphology of the continuum in the final regridded maps, which could therefore be unreliable.

OTF map "striping".

Drift-related behaviour described in the preceding paragraphs may lead the user to wonder if it might manifest as kind of "striping" in affected HIFI maps, similar to the elongated artefacts along the scanning direction of some maps obtained at ground-based telescopes. In ground-based maps, the striping mostly arises from the imperfect correction of the receiver instability due to the atmospheric contribution of the overall signal. The common practice is then to take maps at orthogonal scanning directions, and combine them with more or less sophisticated algorithm (e.g. the PLAIT technique developed by Emerson & Graeve 1988). For HIFI, such a contribution does not exist, so that it is expected that the only residual spurious spatial structures would arise from the imperfect cancellation (in the ON−OFF calibration) of instabilities due to the receiver drifts, such as standing waves or other band-pass distortion, e.g. at the IF edges of diplexer bands (Section 5.3.1). There is also a spatial "zig-zagging" of OTF map scan legs (discussed below) that may be present and might contribute to the appearance of a kind of alternating signal pattern in the map, but this is unrelated to instrumental drift effects.

In order to quantify the striping effect, several cases both in SIS and HEB bands were examined, using maps taken along various scanning directions. The net result is that line peak intensity, or integrated line intensity maps more generally, are not noticeably affected by striping when a proper baseline correction is performed. An example is shown in Figure 5.23, where the baseline drift and standing waves are still significant compared to the observed strength of the C^+ line. In this example, the data are from SPG v14.1 and an ESW correction has already been applied in the pipeline. To improve on the pipeline result, the user must either experiment in HIPE with finding a new ESW solution interactively (with manual line masking in `hebCorrection` – Section 6.4.1), or remove the artefacts and base-

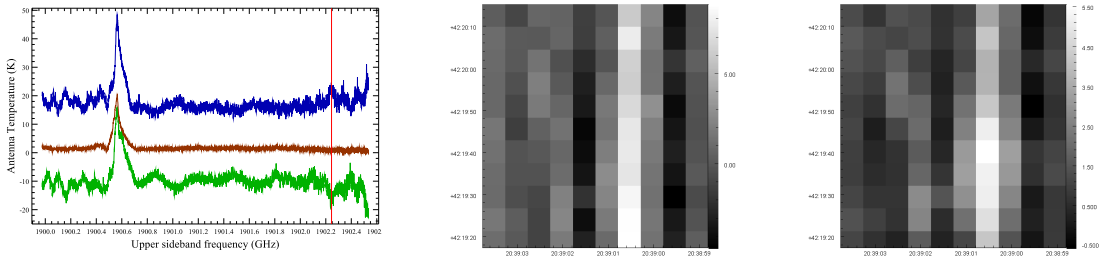

Figure 5.23: OTF map of C^+ emission from DR21 (obsid 1342199677). On the left are spectra with maximum positive and negative drift around the average spectrum (brown), also exhibiting residual electrical and optical standing waves. In the middle is shown the "striping" or "jailbarring" in the baseline due to the drift, which also heavily affects the signal distribution of observed C^+ as shown in the map integrated over the emission line on the right. Both maps are very similar since the average emission line intensity is not much greater than the scale of the baseline drift.

line offsets with a combination of `fitHifiFringe` and `fitBaseline` (Section 6.4.2). The drift present in this observation cannot be separated from source continuum emission, and even though the average spectrum may indicate some continuum flux above zero, the spatial distribution of the continuum is lost in the striping. In those circumstances, any continuum information has to be sacrificed in order to treat the baselines for a reliable line intensity distribution, as shown in Figure 5.24.

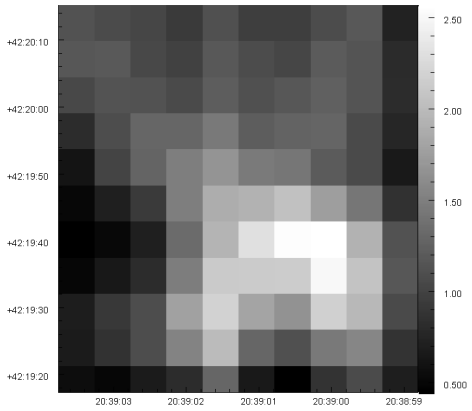

Figure 5.24: Same C^+ integrated map as shown on the bottom right in Fig. 5.23, after baseline corrections using `fitBaseline`

From a general perspective, the ICC discouraged the OTF mode if the science goal required an accurate recovery of the continuum spatial emission. The alternative (if other modes were unfeasible to the Observer due to efficiency constraints or source structure) would be to perform OTF maps along two perpendicular scanning directions. Aside from testing by the ICC, this option was never used, and no tools have been provided in HIPE specifically for combining orthogonal maps (e.g., the PLAIT algorithm, Emerson & Graeve 1988). However a simple tool is available to combine any set of Level 2 products defined by the user (`mergeHtps`), as the input to the HIFI it doGridding algorithm to grid

the signal into pixels according to specified map dimensions and WCS weighting, and create a single output spectral cube.

OTF map "zig-zag".

Early during validation tests of the OTF mapping modes, the ICC found that the astrometry of the map points were measurably shifted from the expected positions on the sky, from one scan line to the next in an alternating fashion. The effects appear as a kind of alternating overshoot and undershoot of the expected map point positions starting in each scan leg, and could be most easily detected in maps in the high frequency bands and for sources with strong spectral emission gradients, such as C^+ from the Orion Bar. The size of the alternating offsets were of the order of a beam size. An example is shown in Figure 5.25. During most of the remaining cold mission the anomaly could not be tracked down as an error in the AOT logic, or a timing error internal to HIFI or in the commanded interactions between instrument and telescope. Neither could the problem be attributed to an issue in the HIFI pipeline, with processing the pointing history product and assigning attitudes to the datasets. In order to compensate the undesirable edge effects and meet the Observer's requested map dimensions, an additional readout was added to each scan line. Some cost to observing time was obviously added as well, and Observers with OTF maps in their programmes were given the option to accept the additional time if it could be afforded, or else make adjustments to the size of the map or select a lower sampling (e.g., Nyquist to Half-Beam).

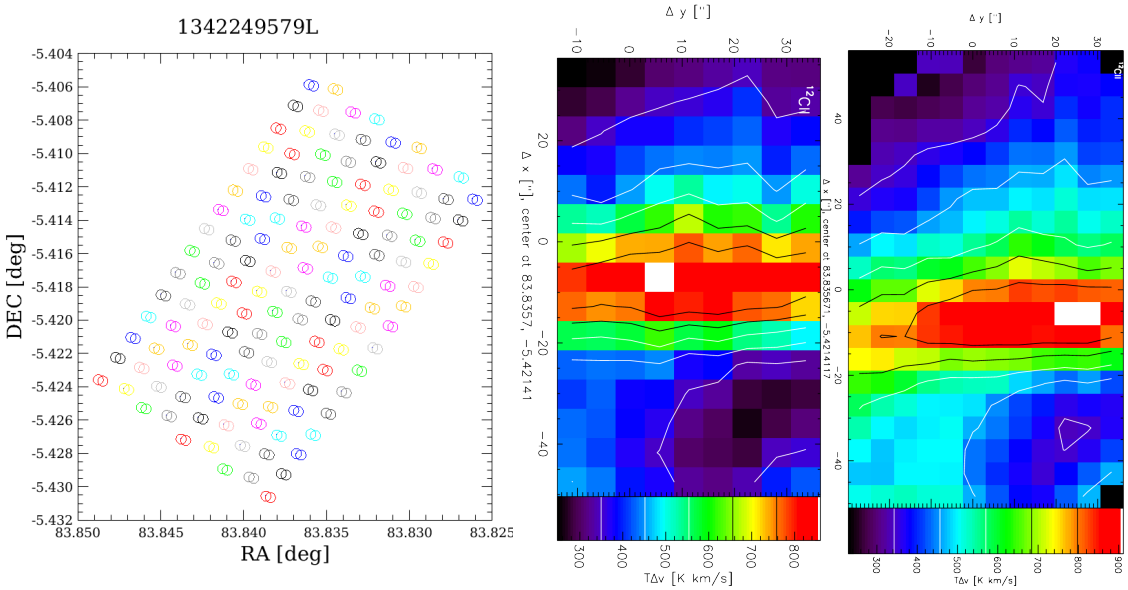

Figure 5.25: Example of OTF pointing "zig-zag", for obsid 1342249579, a map of C^+ from the Orion Bar. The left plot shows the pointing pattern of the ON-target points, taken from the attitude history at that time (circa SPG 8.0). The spatial offsets between neighbouring scan lines are evident. The centre image shows the effects of the pointing on the line intensity map (rotated for clarity). The right image shows the results of an attempt to apply ad hoc timing offsets to the pointing history product for this observation, on the assumption that either a timing error has occurred somewhere between execution of the map and processing of the telemetry, or that it can be treated that way. The result is closer to the expected intensity distribution, but the zig-zag noise was not completely removed, and the ad hoc timing offsets were not suited to general application.

The OTF zig-zag has ultimately been revealed as being caused by anomalous behaviour of the telescope pointing mode `line-scan-with-off` employed by the OTF mapping mode with periodic position switching to a reference sky position. The problem was revealed with the assistance of the PACS ICC during the development of an improved pointing reconstruction algorithm, which uses a more sophisticated approach to measuring the drift on the 4 *Herschel* gyroscopes, yielding a more accurate characterisation of the short-term attitude drift or "jitter" of the telescope. Applied to HIFI OTF maps exhibiting zig-zag, the new pointing reconstruction reveals an unexpected component of drift away from the intended slew path, both spatially and temporally. The path deviations are as much as 3 arcsec (much larger than the telescope APE an RPE, Sánchez-Portal et al. 2014), and appear to occur following a slew of the telescope from the OFF sky reference position, with more of a "banking" movement before converging to the intended path, rather than an undershoot or overshoot of the scan leg's initial intended readout position; see Figure 5.26. Thus the term "zig-zag" is not entirely accurate to describe the erroneous pointing pattern.

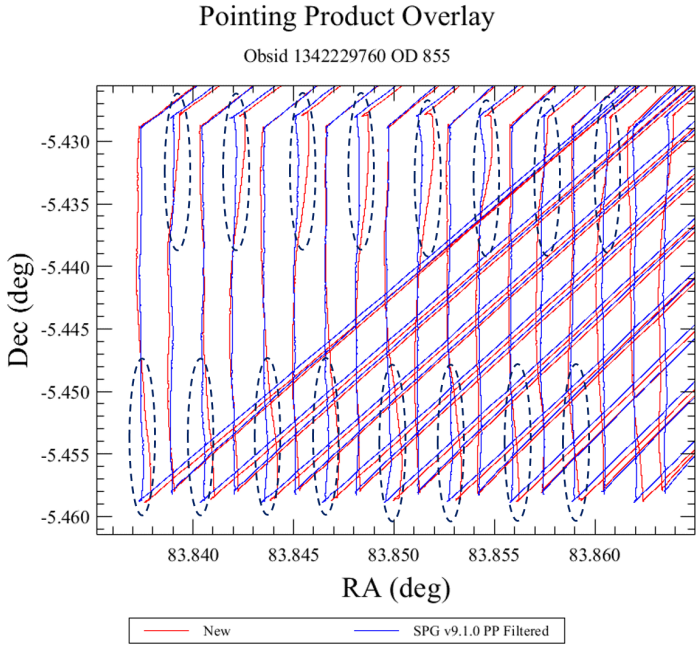

Figure 5.26: A comparison of the continuous slew path of an HIFI OTF map derived standard pointing products (circa SPG v9.1) in blue, and gyro-based pointing reconstruction in red (SPG 13.1 and after). The dashed circles highlight where the gyro-based pointing indicates significant deviations from the expected path (here $\approx 2''$ radially), following slews from the OFF sky reference position.

While the erroneous pointing during the transition from deceleration to acquisition of scan path is strictly on the side of telescope's attitude control system, the HIFI ICC has spent significant effort to determine the scale of the errors and dependencies on slew angle and distance to the OFF position, the slew speed and line scan rate, etc., as well as searching for indications of the problem in other Observing Modes. Summarising this status, in regard final products in the HSA:

- Only OTF maps utilising a position switch (i.e., the `HifiMappingModeOTF`, `HifiMappingModeOTFLoadChop` and `HifiMappingModeOTFFreqSwitch` modes)

may exhibit pointing pattern noise, arising from anomalous deviations of the intended slew path. The `line-scan-with-off` pointing mode errors were not solved before the end of the mission.

- OTF maps most strongly affected would be those taken in the bands with the smallest beams, namely bands 6 and 7, and on science targets which are compact or have structured emission.

- The so-called "gyro-based" pointing reconstruction implemented in SPG v13.1 (and further refined in v14.1), with a more accurate history of the short-term drift component, significantly mitigates pointing pattern noise in the convolved spectral cubes, by virtue of more accurate attitude-based signal weighting during cube construction in the HIFI pipeline. Figure 5.27 shows an example of the improvement, using an early test version of the gyro-based reconstruction method. In fact these results were a significant motivating factor for the HIFI ICC to apply the gyro-based pointing in the SPG to every possible observations.

- However, as discussed further in Section 5.6, not all observations could benefit from the gyro-based pointing, since the reliability of this method relies heavily on the quality of the fitted solutions to gyro an star tracker data. The HIFI pipeline discriminates low-quality gyro-based pointing, and defers to the former "simple" pointing history which is not necessarily more realistic than low-quality gyro-based pointing.

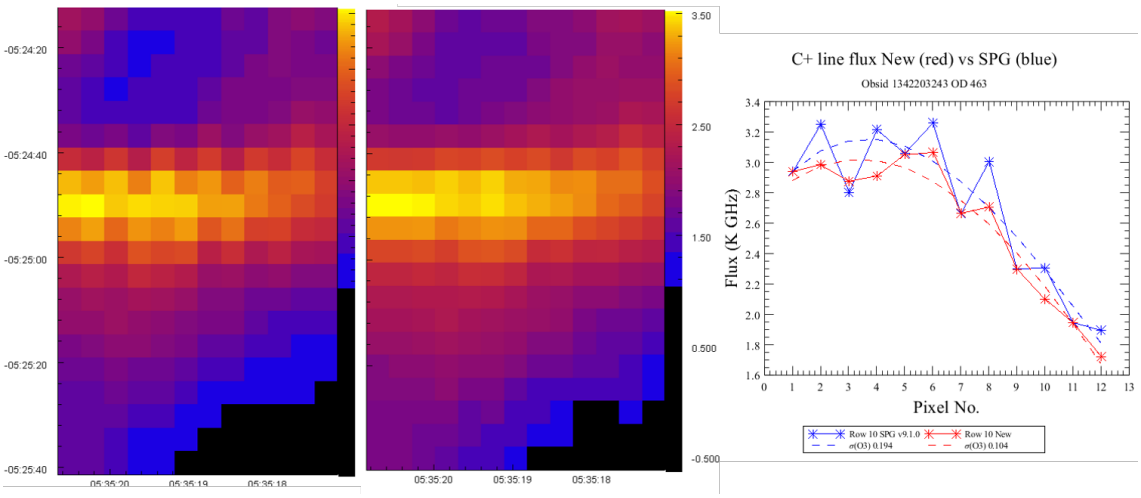

Figure 5.27: A HIFI OTF map of C^+ in the Orion Bar PDR. On the left is the integrated line intensity map from data products using the standard pointing products, circa SPG v9.1 In the middle are the same data, with astrometry from a test version of the gyro-based pointing products, now a part of the SPG since v13.1. On the right is a comparison of integrated C^+ line intensities along row 10 (as an example), in the two versions of the spectral cubes. The blue points are based on standard pointing products, red points are from a development version of the gyro-based pointing. The anomalous pointing pattern noise is clearly lower in the data using a gyro-based pointing product.

Spectral Scanning Modes.

The noise values for the Spectral Scans shown in Figure 5.20 have been measured from the dual sideband (non-deconvolved) H+V averaged spectra at the reference frequency where noise is predicted in HSpot, and have thus been appropriately scaled by the factor $\sqrt{2 \times \mathrm{redundancy}}$. Note that this factor scales to the SSB noise when the gains in each sideband are equal. When the gains deviate from unity, the deconvolution algorithm in the HIFI pipeline should apply these as well and the same scaling should apply. Sideband gain calibrations are summarised in Section 5.8.1.

In the pipeline, the `mkRms` task estimates the noise at the reference frequency in the SSB, H+V averaged spectra, and populates metadata with the estimated values. Numbers from the DSB single polarisation data for the Spectral Scans are anyway plotted in Figure 5.20 to show the reasonably good agreement between predicted and measured noise before deconvolution, to stay clear of potential problems with automated masking of these features to measure the baseline RMS.

While the SSB RMS values are found to be in reasonably good agreement with expectations, some cases are 1.5 to 2 times higher than the HSpot predictions. The worst cases are, not surprisingly, from Spectral Scans using Frequency Switching, and probably related to LO power which has not completely stabilised at each of the throw frequencies. Spectral Scans with Frequency Switching are in case very scarse in the HSA, and have their own peculiarity when it comes to deconvolution, see Section 6.4.2.

Observations using the Load Chop mode, by contrast, have noise levels well below the predictions, on average. The DBS modes are on average above the predictions (somewhat surprisingly). These could be explained by how much or how well interactive "pre-cleaning" of artefacts in the input Level 2 products was carried out before deconvolution, using software (circa HIPE v6.0) at the time of the analysis. Nearly all Spectral Scan observations in the HSA have now been interactively inspected and the flag tables updated for regions of spurious response, and all Spectral Scans in Bands 6 and 7 have been treated for ESWs.

It is clear that final sensitivity is enhanced following sideband deconvolution. Lines that are undetectable in individual double sideband spectra may become visible in the cleaned and deconvolved spectrum. An example of the value of interactively cleaning-up spectra prior to deconvolution is shown in Figure 5.28. Thus the noise comparisons in Figure 5.20 are on the pessimistic side.

Noise performances measured in the HIFI pipeline from final products for each observation in the HSA (SPG v14.1) indicate very good compliances with expectations. In Spectral Scans using the DBS mode, for example, accounting for over 60% of all Spectral Scans with Level 2.5 products in the HSA, 80.0% of the observations have noise levels within the threshold ratio of 1.4 of the measured/predicted RMS noise (SSB, H and V combined). The 20% which are not in compliance (33 total observations) are shown Figure 5.29. In many cases, the high measured noise levels are due to inadequate isolation by the `mkRms` routine of a clean region of baseline around the reference frequency. All four of the highest outliers in Figure 5.29 are surveys of the Orion KL and the Galactic Centre source SgrB2, which are spectrally rich and broad-lined. Thus the plots in Figure 5.29 give a pessimistic view of performances. Other cases do exhibit some residual optical or electrical standing waves, representing insufficient calibration to meet expected performances. Adjusting for the broad-lined or spectrally complex cases, the portion of observations with "true" noise performances issues, i.e. adversely affected by poor baseline quality, is probably lower than 10%. Data quality for these cases will be improved with interactively cleaning of the pipeline products, and made available to the end user as Highly-Processed Data Products (Teyssier 2017a – see also Section 6.2.3).

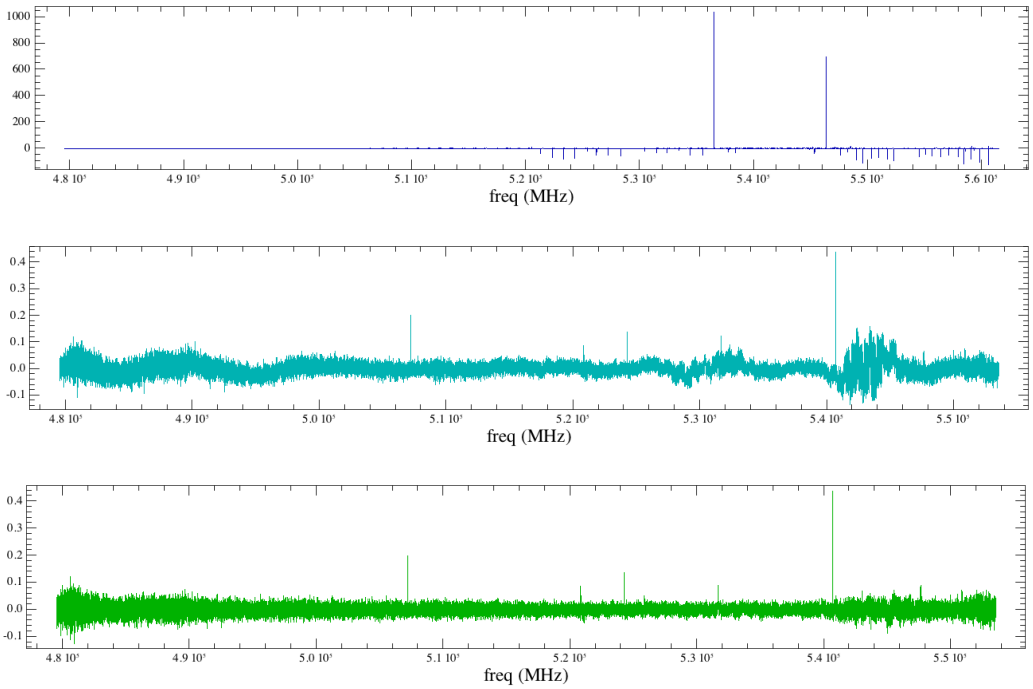

Figure 5.28: Band 1a spectral scan in various states of cleaning. The top plot shows the raw output from deconvolution. The effect of the spurs is obvious, as they echo throughout the solution. The centre plot has the spurs flagged-out, and the solution is much better, though issues with baselines are apparent. After applying a simple baseline removal routine to correct for residual drift and standing waves, the result is further improved as shown by the spectrum on the bottom. The centre plot is representative of the pipeline products one will get when baseline drifts are still present after bandpass calibration. The lower plot is representative of what can be expected from Highly-Processed Data Products offered for Spectral Scans – see text for details.

A similar check on the SPG v14.1 statistics of the FastDBS mode shows lower performances, where 35% are found to be above the `rmsNoise` threshold of 1.4. Again many cases can be attributed to line masking problems, but adjusting for this, the number of observations with legitimate baseline artefact issues is roughly twice that of the DBS mode observations. Fast-chopping has been recommended especially in diplexer Bands 3, 4, 6, and 7 to help offset higher drift rates, where fast-chopping anyway should be more effective at removing standing waves. However we find no convincing correlation between the measured total noise and the drift noise contribution, nor especially with LO frequency or system temperatures at the reference frequency. In relation to the performances of the Point and Mapping AOTs using DBS and FastDBS, these are part of further study for publication.

Statistics for the Load Chop mode shows that 7 out of 24 observations (for which a comparison to the HSpot predictions is possible) are above the measured/predicted ratio of 1.4. However, 6 of these cases belong to Spectral Scans of the source SgrA with many broad and weak lines.

Summarising the Spectral Scan observational results:

- Most Spectral Scans observations have nominal and often quite good noise performances. Cases

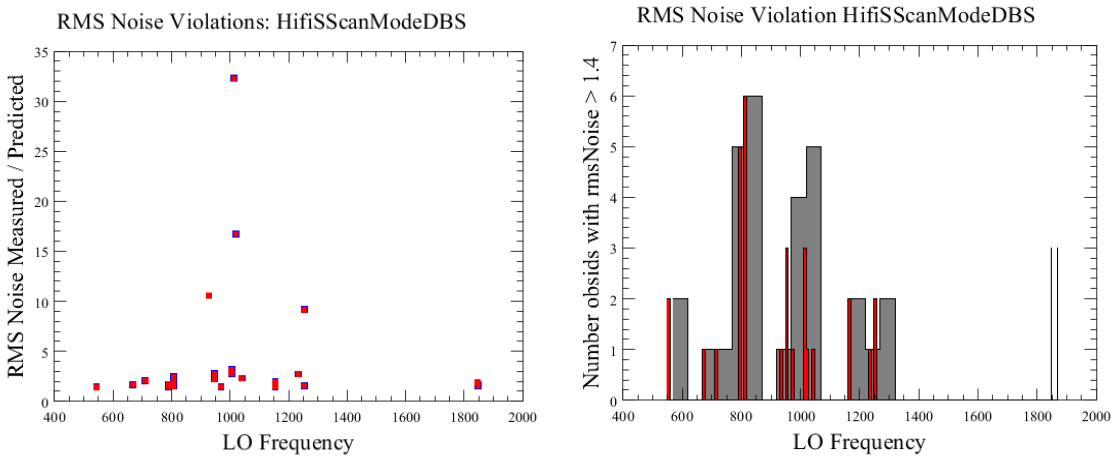

Figure 5.29: Noise performance statistics for DBS mode Spectral Scan observations. The plots show the statistics for observations which have RMS levels at the reference frequency higher than predicted by >40% (a threshold used in the HIFI pipeline to flag noisy cases). The violations are 20.0% of the total number of observations for which HSpot predictions are available for this mode. The highest outlying cases are complex and broad-lined sources, which impedes automated line masking and leads to erroneously high baseline noise values.

which do not meet expectation are often explained by either (a) inadequate automated (pipeline) isolation of a spectrally featureless region of the baseline around the reference frequency, in broad-lined or complex spectra, or (b) uncorrected drift and residual optical or electrical standing waves or high levels of spurious response. These latter cases probably make up ∼10% of the data from DBS mode observations in HSA produced in the final SPG v14.1, while FastDBS data have a greater rate of lower than expected performances, ∼20%. Most observations will have been further processed interactively, and made available in the HSA as Highly-Processed Data Products.

- No ghost lines, or symptoms of instabilities or insufficient sampling, have been found in any spectral survey which is appropriately flagged at frequencies where artefacts such as spurs are present. When spectral artefacts and bad baselines are removed completely from the input prior to deconvolution, the deconvolution yields flat baselines with neither ringing, nor any added repetitive noise structures. An exception to this exists in the event of very strong lines, which can lead to artificial negative features, see Section 6.4.2. Those should be dealt with for most cases in the Highly-Processed Data Products for Spectral Scans though.

5.5.4 IF Edge and Standing Wave Effects on Noise Performances

The effects of optical and electrical standing waves and of high noise near IF edges in certain bands have been summarised in Chapter 5, see especially Sections 5.1 and 5.3. Their characterisation and corrections, in the pipeline and with interactive tools, have been critical for accurate representation in the time estimator noise model during observation planning, and for achieving consistent performances in the data products. Below are highlights of how baseline distortions affect performances

from an Observing Mode context.

Performance/Sensitivities at IF Edges

In the diplexer bands 3, 4, 6 and 7, a substantial increase in system temperatures and thus decrease in sensitivity occurs towards the edges of the IF bandpass. For Bands 3 and 4, the IF bandpass span is 4 GHz (4 8 GHz), and the last 500 MHz at either side, between 4.0-4.5 GHz and between 7.5-8.0 GHz, induces a significant increase of baseline noise since the diplexer mechanism introduces extra losses in those locations, and $T_{\rm sys}$ can increase up to 50-100% as compared to the central part of the IF (see Figure 5.1).

Additional drawbacks include degraded stability performance (stronger standing waves and poorer baseline performance). Those stability issues could be overcome by using FastDBS, but not the sensitivity degradation. Thus it has not been so much a lien or caveat on the Observing Modes, but rather a recommendation to Observers from the ICC to avoid placing lines in the last 500 MHz on either end of the IF in Bands 3 and 4, and in the last 250 MHz in Bands 6 and 7, when using modes of the Point or Map AOTs. In cases where two lines have been targeted in the edges of the upper and lower sideband in a single AOR, it was recommended to devise separate AORs for each line despite the additional overhead, since time is not being saved in a single AOR when the noise and standing waves are impeding spectrum quality.

Standing Wave Residuals

The impact of optical standing waves in HIFI data on line and continuum calibrations is described in Section 5.3. The HIFI pipeline arithmetically corrects for standing waves using the sky reference measurements obtained during observing, only on an individual observation basis. In the DBS modes available with all three AOTs, the sky references come automatically by definition of the switching scheme, at a rate depending on stability parameters and the user's AOR setup, employing telescope nodding plus chopping of the internal M3 mirror to fixed 3-arcminute throws on the sky (Section 3.3.1). The sky measurements are otherwise referenced to a telescope position switch at Observer-selected coordinates within 2 degrees of the main target. If sky position switch measurements were skipped (they were optional but highly recommended with the Frequency Switch and Load Chop modes), then no standing wave correction is applied. Otherwise, in the ideal case, standing waves are effectively removed by action of the `doRefSubtract` step in the Level 1 pipeline (Section 6.1). As mentioned numerous times above, real instrument behaviour involves system drift which is not always well-matched with adopted characteristic Allan times at all frequencies in all bands, resulting in standing wave residuals after standard pipeline processing.

Generally, the beam-switched modes tend to be less susceptible to drift and residual ripples since the ON and OFF positions have reversed optical paths when switched (thus standing waves with opposite sign) so that the ON and the OFF datasets are basically averaged on an individual scan basis. This also makes AORs using DBS, especially at the faster chop rates, more ideal to mitigate drift if the LOs are not completely thermalised after initial switch-on and warm-up. This is most relevant in Bands 6 and 7 where the LOs have very long times to reach thermal stability (Section 5.3.1).

Conversely, total power mode observations employing a telescope Position Switch can be expected to exhibit residual standing waves, since telescope slews are generally long compared to applicable Allan times (Section 5.2). The pipeline must perform some appropriate temporal interpolations on a

baseline of sky reference measurements for multiple scans. LO bands 6 and 7 with the shortest Allan times are thus where residual baseline drift and standing waves are not unusual in some fraction of the calibrated Level 2 spectral datasets. The Spectral Scan AOT and OTF mapping modes may be most susceptible. For both chopped and total power modes, astronomical sources with continuum emission may be expected to show ripples in the Level 2 spectra that multiply between the source and whatever internal residuals exist (Rodríguez-Coira & Teyssier 2015). The impact of the standing waves on science is obviously less and may even be ignored in some cases of sources with very narrow lines and where surrounding baseline or continuum levels are unimportant to the user's science goals.

Baseline drift and standing waves or their residuals after calibration can be mitigated in interactive analysis, using certain tools in HIPE other scientific analysis software, as described in Sections 6.4.1 and 6.4.2.

5.6 Pointing Performance and Impact on the HIFI data

The pointing performance of the *Herschel* telescope and the accuracy of the reconstructed pointing history that is applied to HIFI data in the pipeline can affect scientific interpretations, depending on the mode of observing and frequency-dependent beam size. While HIFI's mixer beams range from $\simeq 11''$ - $43''$ HPBW (see Section 5.9), spatial information can be extracted from high S/N observations on angular scales of a fraction of the beam sizes, which are in turn comparable to the fundamental pointing performances of the telescope statistically represented by the Absolute Pointing Error (APE), Relative Pointing Error (RPE), and Spatial Relative Pointing Error (SRPE). These quantities essentially describe the accuracies of the point-to-point"blind" slew (as HIFI does not employ a pointing peak up), the short-term jitter around a fixed requested position, and precision offsetting (see, e.g,, Section 2.4.4 of the Herschel Observers Manual). The accuracy with which the pointing history is reconstructed is the Absolute Measurement Error (AME). All of these quantities have time dependencies across the *Herschel* mission. The APE, RPE, and SRPE have been measured from dedicated observations of suitable stellar calibration sources with the PACS instrument, treating them as following a Gaussian (random) statistical distribution. Each also has different impacts on the observations from different AOTs and observing modes. Therefore an understanding of how the science can be affected by the bulk performances and potential systematics is beneficial to the use of HIFI observations from the HSA.

5.6.1 Average beam coupling losses

All HIFI observations of compact or semi-extended (non-beam filling) sources will experience beam coupling losses under the influence of random and systematic offsets between the intended and actual pointing. We consider the APE, RPE, and SRPE to be random, while a bias or residual error in the SIAM (Appendix A) is a systematic effect. The theoretical average coupling loss $\langle\eta\rangle$ (an expectation value) in the presence of random (σ_e) and systematic (θ_0) pointing errors as function of σ_b (related to the HPBW θ_b) is given by

$$\langle\eta\rangle = \frac{\sigma_b^2}{\sigma_b^2+\sigma_e^2}\ \exp\Big(-\frac{\theta_0^2}{2\sigma_b^2}\Big(1-\frac{\sigma_e^2}{\sigma_b^2+\sigma_e^2}\Big)\Big). \tag{5.1}$$

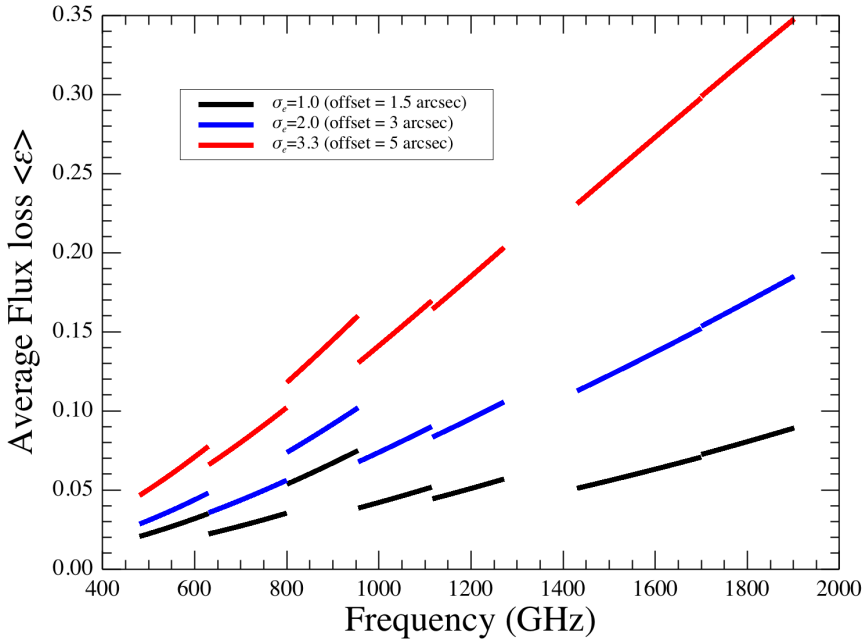

Figure 5.30: Average beam coupling losses at different pointing offsets.

For Point and Spectral Scan AOT observations the random term σ_e is related to the bulk random pointing errors as $\sigma_e = (\mathrm{APE}^2 + \mathrm{RPE}^2)^{1/2}/1.51$, while the SRPE should be included for DBS Raster Map observations. The nominal RPE is less than $0''.3$ over a 60-second period, and it is most meaningful as an RMS variation of signal in multiple readouts at a fixed position or while scanning along an OTF mapping line. The beam-related uncertainty parameter is $\sigma_b = \theta_b/\{2[2\ln(2)]^{1/2}\}$. For all AOTs the systematic term consists of a calibration error in the SIAM θ_{SIAM} and the intrinsic mis-coalignment of the H and V mixer beams θ_{HV}, i.e., $\theta_0 = |\theta_{\mathrm{SIAM}} + \theta_{\mathrm{HV}}|$ (note that they are not uncorrelated).

To give a sense of the expected average impact on individual observations (single fixed pointing integrations) by the random and systematic errors described above, Figure 5.30 shows the variations of fractional coupling losses $\langle\varepsilon\rangle = 1 - \langle\eta\rangle$ over HIFI's range of frequencies using Equation (5.1). For the systematic term it is assumed that the SIAM has a small residual bias error of $1''.5$ in all mixer bands, which is indicated in a mission-wide study of HIFI calibration monitoring sources (Appendix A and Figure 5.46). The H and V polarisation co-alignment offsets are taken from the Focal Plane Geometry measurements, and shown as Δ in Figure 5.44. For the random pointing error, offsets of $1''.5$, $3''.0$, and $5''.0$ are considered. These are approximate multiples of the telescope APE. The HIFI beam at each frequency is also approximated to have a Gaussian profile following $\theta_b = 2.1434 \times 10^4/\nu$ where ν is the LO frequency in GHz.

Figure 5.30 should be interpreted to mean that observations which are offset by less than the 68 percentile of the telescope APE will experience a coupling loss (i.e., average line intensity reduction on a compact source) of less than 10% at any frequency. Higher offsets obviously induce greater losses commensurate with the size of the beams, decreasing in HPBW (θ_b) with LO frequency. Users of HIFI data from the HSA can estimate the effects more specific to the observation using the average offset between the expected and reconstructed pointing. Both of these items are contained in observation

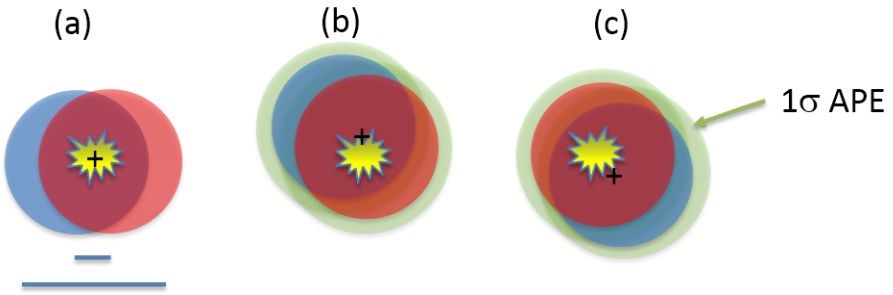

Figure 5.31: Schematic illustration of the H and V beams to scale for Band 3a (a) ideally centred on a compact astronomical source of line emission at the "synthetic" reference indicated by "+", (b) offset by 2″.0 at a random telescope roll angle ϕ, and (c) at the same offset at roll angle $\phi + 180°$. The 6″.3 mis-coalignment between H and V mixer beams and the ≈26″.5 HPBW are indicated by the two horizontal lines. Light green shading in (b) and (c) represents a 1-σ radial uncertainty of 1″.5 from the telescope APE.

context metadata [raNom,decNom] (the requested pointing) and [ra,dec] (the reconstructed pointing averaged over all Level 2 science datasets). If expressed as a radial offset Δr, then Equation 5.1 can be rewritten with $\sigma_b = 9.1022 \times 10^3/\nu$ and $\sigma_e = \Delta r/1.51$ where ν is in GHz and Δr is in arcseconds. The systematic (SIAM) offset θ_0 can be adopted to be 1″.5. Given a value of Δr for an observation at frequency ν, the user can also simply read the coupling loss from Figure 5.30.

5.6.2 Pointing-induced H and V line intensity imbalances

As mentioned above and discussed in Section 5.9.1, there are strong indications from analysis of 3×3 raster maps of a number of calibration monitoring sources that the final SIAM contains an average bias of $(\Delta\theta_Y, \Delta\theta_Z) = (-1.0 \pm 0.8, -1.6 \pm 1.1)$ arcseconds on the telescope focal plane axes [Y,Z] as projected onto the sky (Figure 5.46). Since the H and V polarization mixer beams are also not perfectly co-aligned (see Figure 5.44 and Table 4 in Roelfsema et al. 2012), the SIAM biases can cause imbalances between spectral line strengths measured through the two mixer beams, depending on spatial structure of the source. This situation is illustrated in Figure 5.31, in which the error-free situation shown in (a) has the synthetic reference aperture centred on the intended target, while in (b) there is a 2″.0 offset at arbitrary telescope position angle ϕ from the intended position and in (c) at a different position angle. The point regarding (b) and (c) is that one of the same mixer beam will be favoured by an offset regardless of the position angle since the telescope focal plane with a fixed geometry of the beam apertures will roll *around* the intended position.

The quantitative impact of these SIAM biases has been calculated (Jellema 2015), treating each mixer band separately for their specific SIAM errors, beam parameters, co-alignment errors, and an assumed APE of 1″.6. The results are given in Table 5.5, where the expected H/V imbalance is given by $\langle \rho_{HV} \rangle$.

The calculation results in Table 5.5 do not take into account differences in coupling coefficients for the two beams, i.e., they are suited to measurements of HIFI observations on the T_A^* scale. Indeed most of the results by band are consistent with a survey of the Level 2 spectra from all Point AOT observations in the HSA (Van Dyk et al. 2016), in which an automated line detection algorithm has been applied to deduce H/V line intensity ratios. Figure 5.31 shows an example for Band 3a and 3b. A great deal of filtering of the data for errant results has been applied, but the median ratio $\langle \rho_{HV} \rangle = 0.95$ is in good

Band	θ_b (arcsec)	$\Delta\theta_0^H$ (arcsec)	$\Delta\theta_0^V$ (arcsec)	$\langle\eta_H\rangle$	$\langle\eta_V\rangle$	$\langle\rho_{HV}\rangle$	σ_ρ
1	33.0	4.1	3.4	0.95	0.97	0.98	0.03
2	26.6	4.1	1.4	0.93	0.98	0.94	0.03
3	22.1	4.4	2.1	0.89	0.96	0.92	0.07
4	18.9	3.6	0.4	0.89	0.98	0.91	0.05
5	17.3	0.8	3.3	0.97	0.89	1.10	0.06
6	12.4	2.0	2.6	0.90	0.85	1.05	0.03
7	11.2	3.2	2.3	0.76	0.85	0.90	0.04

Table 5.5: Impact of SIAM bias errors on the scientific calibration of H and V line intensities.

agreement with the expected value of 0.92. Note that if beam efficiencies are taken into account and the calculations and measurements are done on the $T_{\rm mb}$ scale, and *assuming that celestial sources are compact*, then $\langle\rho_{HV}\rangle$ increases everywhere except in Band 1. Bands 2 and 3 in particular have $\langle\rho_{HV}\rangle$ predicted to be within 3.0% of unity.

5.6.3 Pointing reconstruction.

The pointing history is calculated on an Observational Day basis and formatted into products for each observation as part of Standard Product Generation at the HSC. Over the time range (in UTC) of every Obsid, an Attitude History File (AHF) and a Pointing Product extracted from the OD-based products are provided as products in the Auxiliary data tree for each observation under the Observation Context. The Pointing Product in particular provides the usable timeline of attitudes in the frame of the spacecraft Attitude Control Axis (ACA) in which the x-axis is the telescope boresight. The `doPointing` task in the HIFI pipeline applies the astrometric rotations to data taken with the appropriate H and V mixers based on the SIAM.

Starting in SPG version 13.0, the calculation of the pointing history incorporates a new method to estimate the drift on the four gyroscopes (in short hand the "gyro-based method") more accurately, leading to a higher fidelity representation of the short-term pointing stability or so-called jitter of the telescope, and is quoted as the RPE. The method was developed and first introduced into HIPE for interactive testing by the PACS-ICC (Feuchtgruber 2012). The RPE is measured from integrations on a target between telescope moves, and is quoted to have a requirement of 0.3 arcsec (1-σ radial) when measured over 60 sec. To HIFI with beam sizes between ≈11 and 43 arcsec HPBW, this quantity may seem inconsequential. However the improved characterisation of the drift on the gyros also provides lower reconstructed attitude errors, cast in terms of the AME, in general. This can be viewed as a reduction to the error of the reconstructed path of the telescope, whether stationary or making short slews or continuous scanning motion during an observation. The importance of this to HIFI was recognised very early particularly for mapping observations in the development of the method through collaborative testing with the PACS-ICC.

The benefits of the reduced AME to HIFI observations has been summarised in a note written for users of HIFI data from the HSA by the HIFI-ICC (Morris 2015), and the reader is referred to that document for a description of results as they apply to data processed in SPG 13.0 compared to SPG 12.0. The effects are also described with the AOT validations and performances in Section 5.5. Here we emphasise three key points about the "legacy" pointing reconstruction in relation to HIFI data:

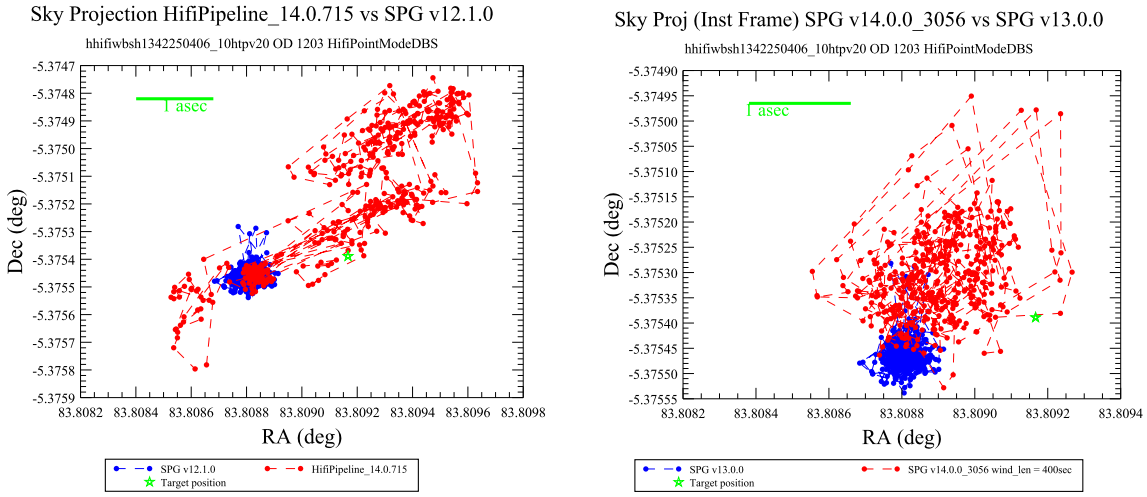

Figure 5.32: Example of an observation for which the gyro-based pointing reconstruction in SPG 14.0 (left, red) indicated very poor stability compared to the "simple" pointing provided in SPG 12.0 (left, blue). However, the quality of the gyro-based pointing has been determined to be very low since the sensor data are affected by guide star quality issues resulting in a very low probability of a good fit to the sensor data on each spacecraft axis, producing a highly unreliable short-term stability component. Thus the attitudes which were applied to these data in the archive processed with SPG 13.0 are rolled back to the simple pointing reconstruction. The gyro-based pointing in SPG 14.0 (right, red) are based on better recognition and treatment of the unruly guide star data, thus the resulting attitudes have higher estimated quality (more representative of the telescope RPE and SRPE) and are applied to the data in the archive.

1. While the gyro-based method frequently reduces the AME and thus improves the detailed pointing history on an observation basis, it suffers some drawbacks. The most significant one relates to the quality of attitude solutions derived from star tracker telemetry, which the gyro-based method is designed to refine. The quality of properties of cataloged guide stars available to the star tracker during an observation translates directly to the quality of the star tracker output attitude solutions, and if the star tracker solution has a lower quality due to one or more guide stars with a quality issue (such as photometric uncertainty or variability, crowded region, etc), then the final pointing reconstruction using the gyro-based method will also have a low quality. This can result in suspiciously large spread around the intended pointing, giving very poor impressions of the relative stability of the telescope. For that reason, a scheme of quality metrics were introduced as metadata into Pointing Products starting in SPG 13.0. The doPointing task in the HIFI Level 0 pipeline processes the Pointing Product with each Obsid, and carries out four additional steps with the quality information, which are:

   (a) Read the quality flag and copy this to the HIFI Observation Context header, for information.

   (b) Raise a suspicious pointing flag which is a simple True/False Boolean, and set to True when the Pointing Product quality flag is below 0.8, and copy this to the Observation Context and HIFI timeline products. When True the flag will also appear in the Quality Summary table produced at the end of the pipeline.

(c) Revert to "simple" pointing reconstruction provided in the Pointing Products in addition to the "gyro-based pointing", when the quality flag is below a cutoff value of 0.4. This basically amounts to reading a different quaternion column in the Pointing Product.

(d) The Pointing Product column which was ultimately applied to the HIFI data is reported as a metadatum in the Observation Context and HIFI timeline products.

The pipeline also compares the intended position as entered by the original observer in HSpot to an average of positions of all Level 2 spectrum datasets filtered down to only the ON position (i.e., excluding sky reference and internal calibration spectra). When the difference is more than 3 times the APE, an "anomalous pointing" flag is raised and reported as a ratio in the Observation Context and in the Quality Summary product. The mission phase-dependent value of the APE is stored in the HIFI calibration tree.

2. The susceptibility of the gyro-based method to guide star quality has been reduced for many observations with improvements to the software for SPG 14.0. As a consequence, a number of observations which had very low quality gyro-based pointing reconstruction in SPG 13.0 have good quality in SPG 14.0 and thus the attitude history for these cases has been switched to the gyro-based reconstruction. The improvements apply only to observations which occurred in the second half of the mission (after OD-880) when downlinked guide star telemetry was sufficiently complete for the procedures to correct for one or more unruly guide stars. An applied example of a case which switched from very poor quality SPG 13.0 to good quality in SPG 14.0 in shown in Figure 5.32. Not all observations in the second half of the mission with poor gyro-based results show sufficient improvement to make the switch from simple pointing. It must be emphasised that the simple pointing reconstruction is not more reliable for observations to which it has been applied (i.e., where the alternative gyro-based results are low quality), just that the misleading relative stability as shown in the left plot of Figure 5.32 has been excluded.

3. Most of the difficult cases for the gyro-based method occur when star tracker interlacing of more than 9 guide stars has been active, and one or more of these has a quality issue. These cases occur most frequently for science targets in the Galactic Plane (noting that the star tracker views a $5^\circ \times 5^\circ$ field along the telescope's $-$x boresight axis). *This science-based selection effect significantly contrasts with how the telescope bulk pointing performances have been estimated.* First, they are derived from PACS observations of point sources selected for pointing calibrations (Nielbock 2010), generally unresolved stellar sources located in well-behaved regions (in terms of guide star fields). A survey of the Pointing Product quality flags shows that there are very few if any problems in pointing reconstruction quality for these observations, ruling out guide star quality issues. Since these observations are the basis of performance studies of the telescope pointing, and the targets have been returned to a number of times, there is a strong weighting towards data from well-behaved parts of the sky, and the results published by Sánchez-Portal et al. (2014) only represent the most ideal performance of the pointing system. Those results are not reflected where a significant amount of *Herschel* science observing time has been invested by all three science instruments. Similarly, the mode of observing used to acquire the pointing calibration data was almost entirely restricted to the PACS chopped point source mode. As such the descriptions by Sánchez-Portal et al. (2014) does not cover some particular behaviour in other observing or pointing modes, for example the anomalous line scanning in HIFI OTF maps (Section 5.5).

The user is advised to regard currently published bulk performance parameters and descriptions to reflect when the telescope pointing system is on its best behaviour, and most suited to PACS point source photometric observations in well behaved regions of the sky; significant deviations in performances and degraded pointing reconstruction quality will occur in regions with sub-nominal guide star qualities, and in other observing modes.

5.7 Spectral Resolution and Frequency Scale Accuracy

The HRS and WBS spectrometers operated on different principals and the calibration requirements were quite distinct. As a digital autocorrelating spectrometer (ACS), the HRS did not require frequency calibration *per se*.

The WBS **did** require frequency calibration, however. Due to ambient temperature changes during operation, the instrumental frequency resolution and range could change. Periodically during observations a COMB spectrum was measured by the WBS and used to calibrate frequency response. Uncertainties arising from this procedure dominated the WBS frequency accuracy. Additionally there were occasional artefacts in the COMB spectra that had to be specially treated.

5.7.1 HRS Frequency Calibration

The HRS bandwidth, channel spacing, and resolution were configured by reallocating the primary instrument resource: digital memory. There were a limited number of lags with which to compute an autocorrelation. Table 2.3 lists the standard HRS operating modes.

As long as the digital electronics remained phase-locked to the master oscillator (and the MO was stable), assigning IF frequencies to the output channels was a matter of correct conversion factors in the control and reduction software.

Only a couple of isolated HRS data corruption events were observed during the mission, leading to improperly calibrated data, either in intensity or frequency. Those data have been accordingly discarded and/or flagged (Appendix B).

5.7.2 WBS Frequency Calibration

COMB spectra were taken after every retuning of the LO, as well as after every measurement on the internal thermal loads. In the case of Spectral Scans this was after every LO frequency change. The COMB cadence was basically driven by the one applying to internal load measurements, which in turn was driven by the corresponding Allan variance times (Tab. 5.1). Typically, these ranged from 20 minutes in most SIS bands, down to 8 minutes in HEB bands.

Figure 5.33 is an example of a calibrated COMB spectrum. Calibration involves first the usual correction of CCD data for dark current and response. Then the harmonic tones, which are known to be separated by 100 MHz, are found along the pixel axis. The tones are unresolved and the WBS spectral response is close to a Gaussian profile, so Gaussians are fitted to the 11 COMB lines per CCD for the four CCDs. The fit provides measures of the spectral resolution, and samples of the mapping from pixel number to IF frequency. A 3^{rd} order polynomial representation of the mapping is estimated from the COMB line fits, a cubic model being expected *a priori* and also found to be best supported by the data.

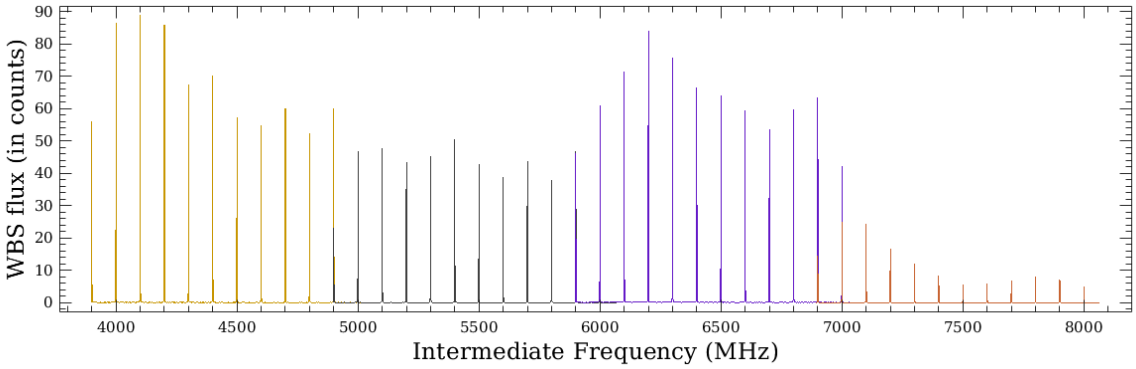

Figure 5.33: A calibrated COMB spectrum from the WBS-V backend, one of 27 in observation 1342249614 (OD 1190). The four colours correspond to the four CCDs in the WBS-V spectrometer. After calibration the COMB tones are at multiples of 100 MHz in IF frequency, by construction. Note that the COMB strength on CCD4 (6900 MHz – 8100 MHz) is less than the others; the laser illuminating the Bragg cell suffered declining power over the mission but never became un-calibratible (Fig. 5.35).

WBS In-Flight Frequency calibration

In-flight tests were conducted to determine the accuracy of the WBS frequency calibration (Avruch 2011). The CO (5-4) line in NGC7023 was observed in Spectral Scan mode such that the line was detected at many positions across the IF as the LO changed. The expected change in IF frequency of the line is the change in LO (with a small adjustment for changing Doppler correction); the actual change in IF frequency was estimated by fitting for a shift between the line profiles at the two LOs (similar to a cross-correlation analysis). The result of these experiments was that the measured shift was in all cases within $\pm$ 100 kHz of the expected shift, see Figure 5.34.

WBS Frequency Calibration Anomalies

As can be seen in Figure 5.33, the COMB power for CCD4 of spectrometer WBS-V suffered a decline over the course of the mission. Calibration tests mentioned above showed that, nevertheless, the COMB calibration remained valid through the end of operations.

An alternative WBS frequency calibration mode for the eventuality of a dead WBS-V COMB (the so-called HRSxWBS, using correlation of HRS and WBS simultaneous spectra to calibrate the WBS) was implemented in the on-board software and in the HIFI pipeline. Fortunately, a switch-over was never required. The pipeline can process test observations made in this mode, but the software module was still under development when it became clearly unnecessary so was left unfinished. The test observations (3.3.4) have an estimated spectral calibration accuracy for WBS-V of $\lesssim$ 500 kHz so have not been released for public usage.

Another occasional anomaly in WBS frequency calibration was due to particles penetrating the instrument and interacting with the CCDs (or possibly cross-talk on signal lines) during COMB integra-

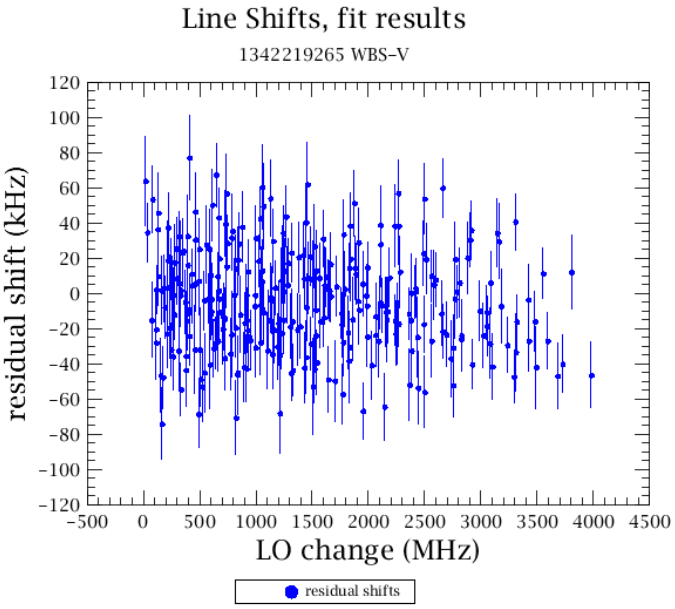

Figure 5.34: Experiments were performed in-flight to test the WBS frequency calibration accuracy. The results of one such test are shown here, where residual shift means the difference between expected and measured spectra line IF frequency shift with a known LO change. In all cases examined an upper bound of $\pm$ 100 kHz on the calibration accuracy held.

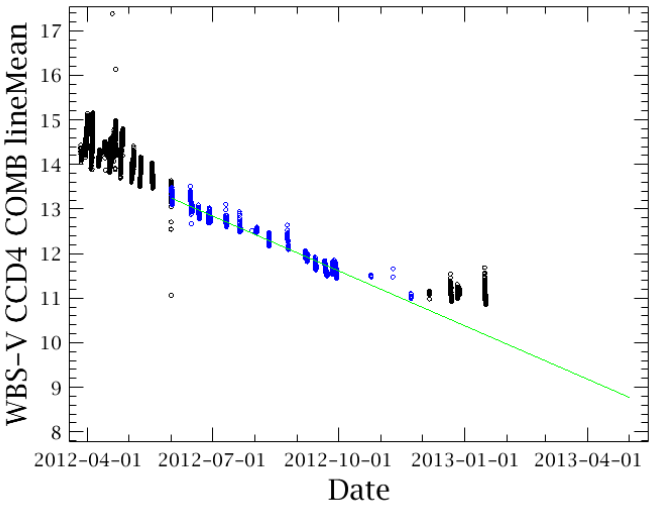

Figure 5.35: The trend of the WBS-V CCD4 COMB power near the end of the mission. A mean level of 8 counts was considered the threshold for abandoning the COMB as primary frequency calibration. We estimated the level might fall to 8.8 by 01 May 2013; in the event the trend flattened and the COMB was usable through the end of normal operations on 29 April 2013 when the cryogenic Helium was exhausted.

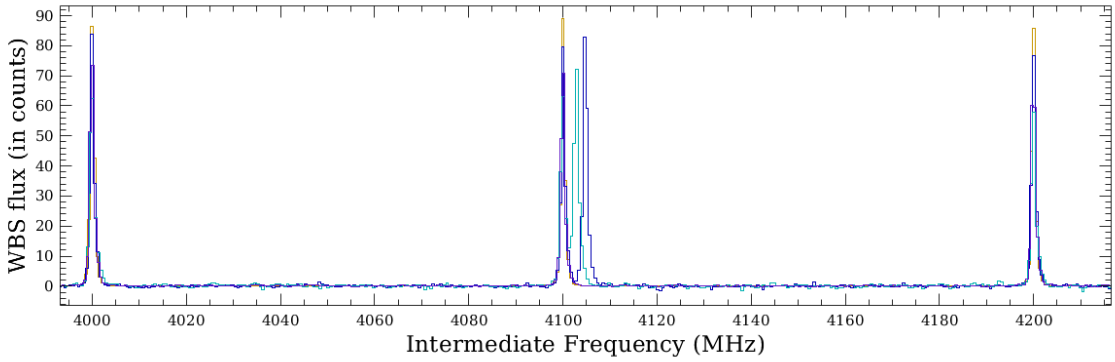

Figure 5.36: An example of a probable charged-particle interaction with a WBS CCD that potentially corrupted the WBS frequency calibration. The two narrow features near 4100 MHz are an artefact in two successive COMB spectra that by chance appeared close to the actual COMB lines. The artefact persisted in time between the COMB spectra; it was present in the intervening science integrations as well. These events were rare (in terms of frequency mis-calibration) and are automatically flagged by the HIFI pipeline.

tions. In figure 5.36 an event occurred that corrupted two COMB spectra taken 69 seconds apart. The anomaly was also present in the spectra taken before and after these COMBs. Despite the similarity, in profile and location, to an actual COMB line, these events could be handled automatically after changes to the pipeline.

5.7.3 Frequency Transformations in Data Processing

Considering the high resolution and stability of the HIFI spectrometers, care must be taken during pipelining not to introduce frequency scale inaccuracies solely due to software. At the highest observing frequencies the frequency resolution is $R \approx 10^7$ and the frequency stability about 1 in $\sim 10^8$. Transforming observed frequencies in the spacecraft frame of reference to other frames such as $\mathrm{LSR_k}$ requires a Doppler correction. Expanding the relativistic Doppler formula in a series ($\beta \equiv v/c \sim 10^{-4}$):

$$\frac{\nu - \nu_0}{\nu_0} = \underbrace{-\beta}_{\sim 10^{-4}} + \underbrace{\frac{1}{2}\beta^2}_{\sim 10^{-8}} - \underbrace{\frac{1}{2}\beta^3}_{\sim 10^{-12}} + \cdots$$

So to preserve instrumental accuracy in Doppler calculations, Lorentz transforms are necessary and the accuracy of the *Herschel* velocity used must be $\lesssim 1$ m/s. Fortunately, *Herschel* orbit reconstruction was more than precise enough.

The pipeline default frames of reference for reduced spectra are $\mathrm{LSR_k}$ for fixed targets, and the source frame for moving targets. More detail on Doppler corrections and frame changes are given in Avruch & Melchior (2011).

For ease of use, the final WBS spectra are rebinned to a regular 0.5 MHz grid, or about 0.45 bins per spectral PSF. The regridding algorithm will change the apparent profile of unresolved lines but preserves total flux.

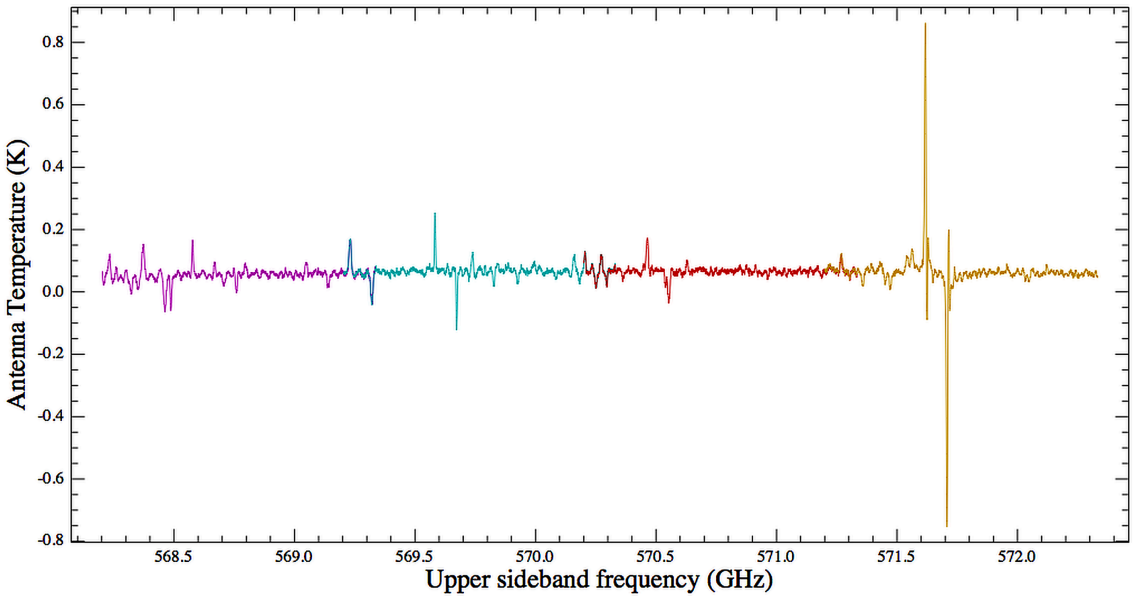

Figure 5.37: Example of a FSW observation in Obsid 1342180473. Note the co-existence of positive and negative features belonging to the respective LO1 and LO2 tunings. Data are shown in an Upper Sideband (USB) scale

5.7.4 Sideband frequency scale and Frequency Switching mode data

Although this is not strictly related to the frequency accuracy of the HIFI data, there is some special care to be taken when interpreting Level 2.5 data taken in Frequency Switching modes. In this mode, indeed, Level 2.5 products will be provided after the *folding* of the spectra is performed (3.3.1, see also the Frequency Switching cookbook). When this transformation is performed, lines will appear differently depending on whether they belong to the USB or the LSB bands. The general shape will be that of one main line bracketed by two satellite lines in the opposite direction, and of about half the intensity. The main line will be positive only if it is calibrated in the frequency scale of the sideband it belongs too – a negative central line will imply it belongs to the other sideband. This is illustrated in Figures 5.37 (unfolded Level 2) and 5.38 (folded Level 2.5).

The general message here is that you can use the Level 2.5 data to figure out what sideband a given line belongs too, a degeneracy that is not immediately resolvable in the other observing mode providing DSB data.

5.8 Intensity calibration accuracy

5.8.1 Sideband ratio calibration

As described in Sections 4.1.1 and 4.6, the mixer sideband gain ratio (thereafter SBR) is a key component of the flux calibration and of its accuracy. Originally, the derivation of the SBR for HIFI

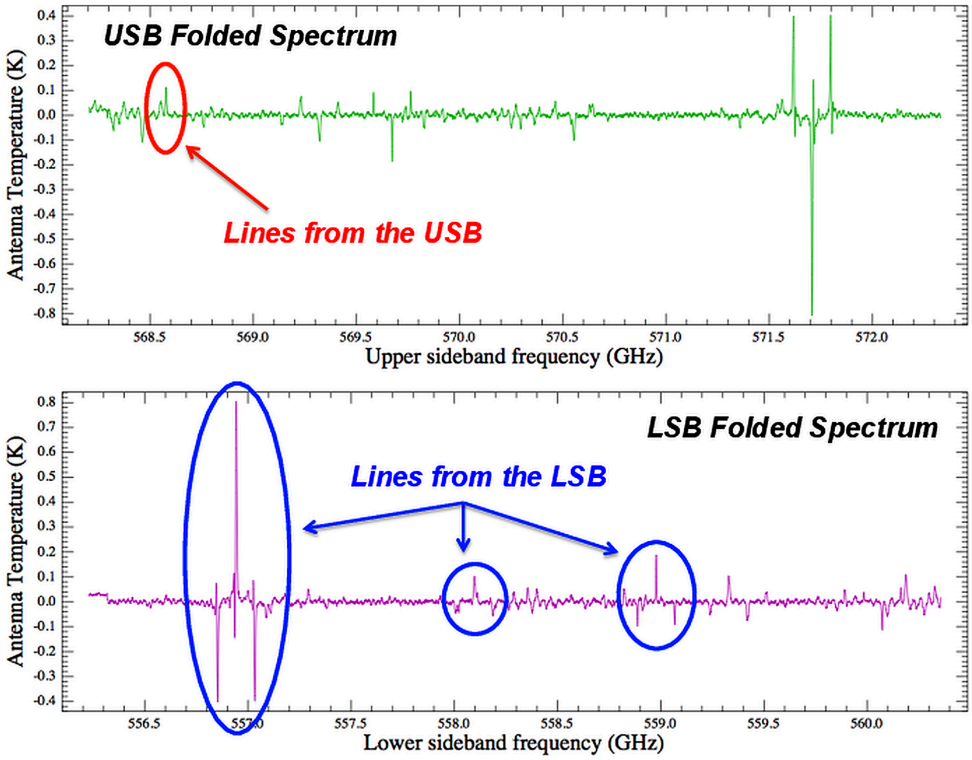

Figure 5.38: Line sideband assignment after the folding. Upper panel: USB scale spectrum. Lower panel: LSB scale spectrum.

was essentially relying on the test campaigns conducted on the instrument prior to the launch. Those laboratory gascell measurements resulted in a set of high accuracy measurement points at isolated frequency tuning, however, it failed to provide a finer granularity picture of how the mixer gain would vary over the entire HIFI tuning range (Higgins 2011, Higgins et al. 2014).

A novel approach was proposed by Kester et al. (2016), that allowed to combine prior information on the SBR from the pre-launch measurement points, to the wealth of redundant information provided by the many Spectral Scans performed on strong sources over the mission. The method, which in essence performs a deconvolution of the double-sideband data, allowed to derive both the SBR on a 1 GHz grid over the entire HIFI frequency range, as well as the accuracy of that parameter, using the power of the Bayesian statistics. Figure 5.39 is taken from Kester et al. (2016) and shows the sideband gain finally used in the HIFI pipeline, together with the associated error bars applying to the overall flux calibration uncertainty described in the next section.

5.8.2 Absolute calibration accuracy

The absolute calibration accuracy is related to the uncertainty model described in Section 4.6 and its propagation into the calibration equation following Ossenkopf (2015) (Section 4.1.1). Since the un-

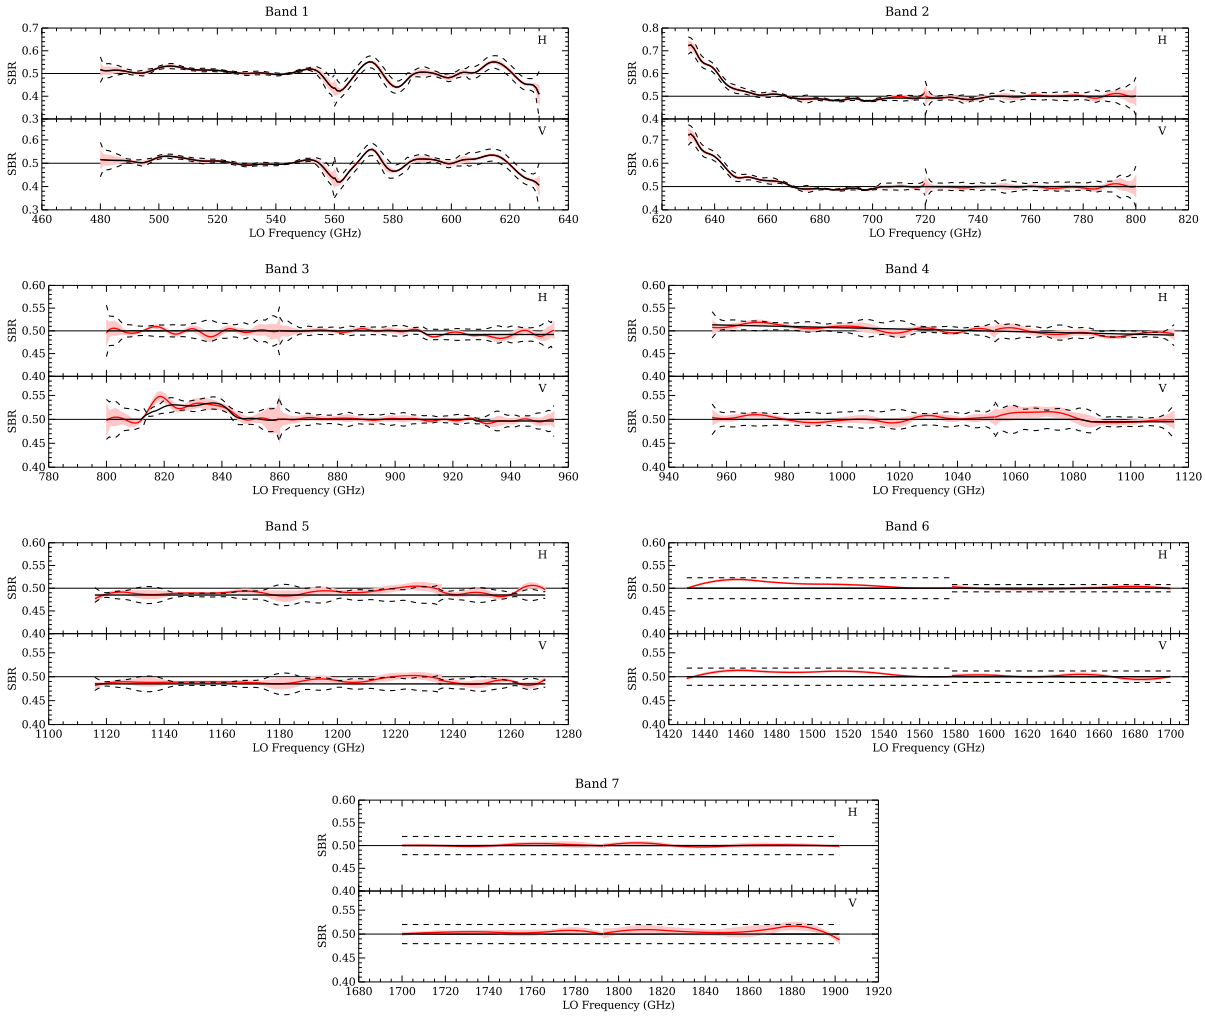

Figure 5.39: Sideband ratio G_{ssb} from the USB (see Section 4.1.1 for conventions) as a function of Local Oscillator frequency for all 7 HIFI bands. The red lines show the result of the Spectral Scan data fit with the 1σ confidence region shown as a pink envelope. The black lines represent the finally adopted values in the HIFI pipeline 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. Taken from Kester et al. (2016).

certainties applying to each of the parameters entering the equation can be LO Frequency-dependent, the total uncertainty will also show a frequency dependency which in some areas has quite some fine granularity.

Table 5.6 summarises the typical uncertainty ranges for each band and each component, as well as the resulting quadratic sum of those, which will apply to the output of the pipeline (Level 2 or Level 2.5), i.e. to data calibrated in the T_A^* scale. These uncertainties should be taken as $1\,\sigma$ root-mean-square (RMS) errors. Fig. 5.40 illustrates the frequency dependence of each of the T_A^* calibration uncertainty components, together with the summed contribution. Since the T_A^* scale is usually not the final intensity scale to be used for scientific interpretation, conversion to other scales implies

Uncertainty component	**Band**						
	1	**2**	**3**	**4**	**5**	**6**	**7**
Mixer sideband gain ratio	0.3–4	0.2–4	0.6–3	0.5-2.3	0.3–2.5	0.8–2.3	2
Internal hot load coupling	0.1	0.4	0.3	0.2	1.7	2	2.3
Internal cold load coupling	0.2	0.2	0.1	0.1	1.7	1.5	1.5
Internal hot load temperature	0.6	0.6	0.6	0.7	0.7	0.7	0.8
Internal cold load temperature	0.1	0.1	0.1	0.1	<0.1	<0.1	<0.1
OSW: internal load	1.4	1.1	0.7	1.2	1.1	0.8	0.8
OSW: diplexer	0	0	0.9	0.9	0	4.2	4.9
Total uncertainty on the T_A^* scale	**1.6–4**	**1.3–3**	**1.5–3**	**2–3**	**3–4**	**5–5.5**	**6**
Beam efficiency	1.1	1.4	1.3	1.3	1	1.2	1.3
Aperture efficiency	1.1	1.4	1.3	1.3	1.1	1.2	1.3
Planetary model	5	5	5	5	5	5	5

Table 5.6: Ranges of absolute intensity calibration uncertainty budget per band, together with the total uncertainties applying to the T_A^* –calibrated data. All uncertainties are given as a 1-σ root mean square error in percentage, apart from the planetary model which should be considered a peak-to-peak uncertainty.

that additional uncertainty components need to be taken into account, starting with the systematic uncertainty of typically 5% (peak-to-peak) applying to the Mars model used to derive the coupling efficiencies.

What this means for the end user is that the total uncertainty to be quoted when using HIFI data is the sum of the uncertainties applying to the T_A^* calibration, to the Mars brightness temperature model (systematic uncertainty of order 5%), and that applying to the respective source coupling efficiencies (again treated as a random error). Uncertainties on beam and aperture efficiencies should not be combined since only one of them has to be involved in the conversion to either $T_{\rm mb}$ or Flux density (Jy) respectively. It should be noted, however, that while random uncertainty on the T_A^* are treated as a 1 σ errors, the uncertainty on the Mars model is a peak-to-peak error, i.e. represent the outer bounds of possibilities of the model.

The T_A^* uncertainty budget and total numbers are actually part of the end products and are provided as a TableDataset in the observation context of each observation (Section 6.1). These are computed based on the uncertainty model stored in the HIFI calibration files (Section 6.5). Those tables are also provided stand-alone as a dedicated expert-curated product in the framework of the Highly-Processed Data Products (Section 6.2).

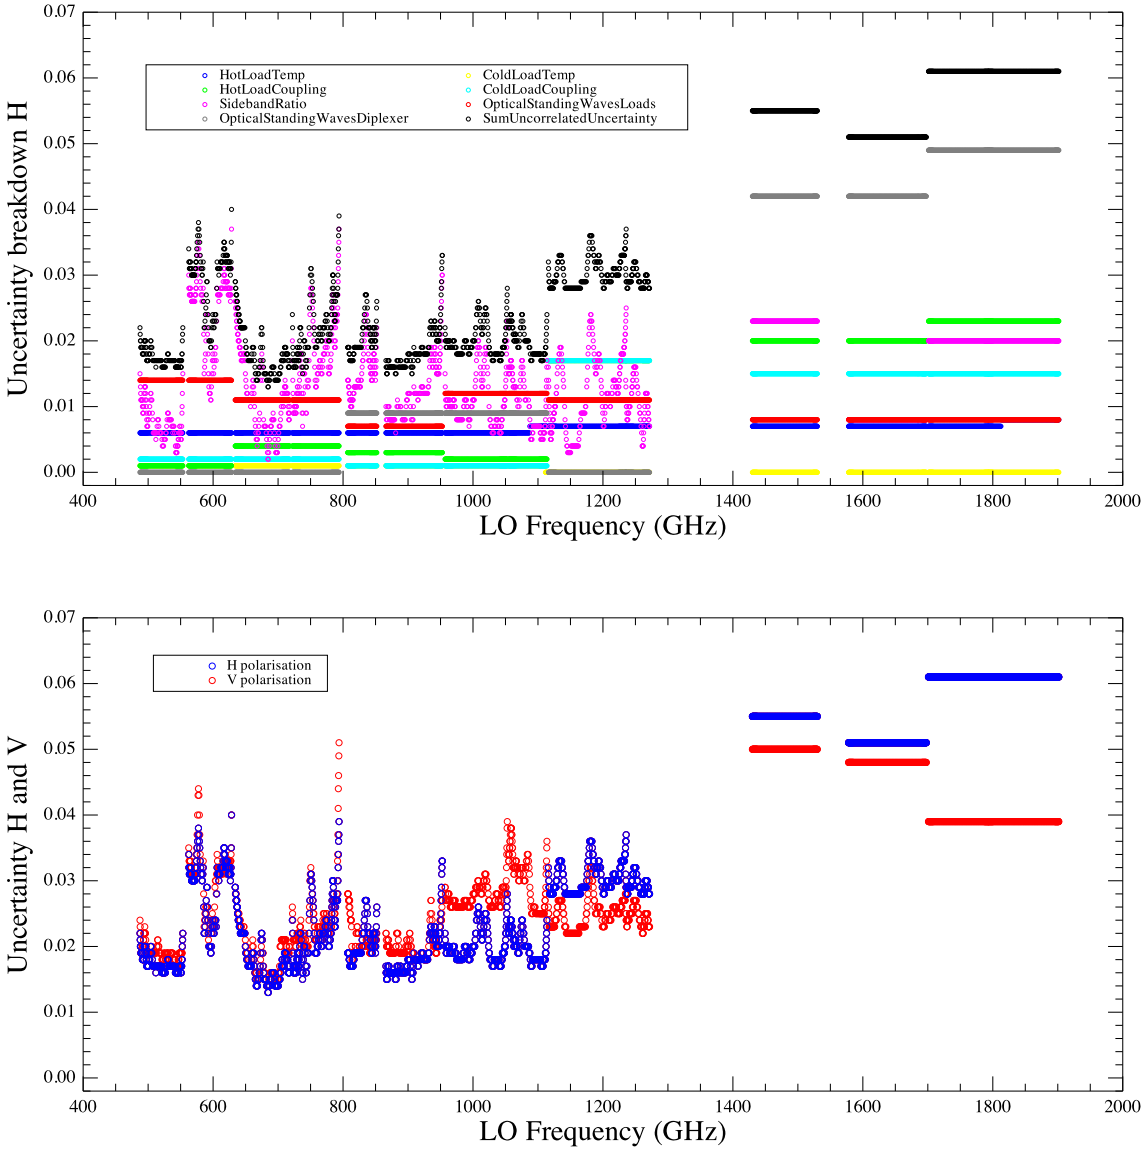

Figure 5.40: Intensity calibration uncertainty breakdown for data calibrated in the T_A^* scale for the H polarisation (upper plot), and sum of the corresponding uncorrelated uncertainties for H and V (lower plot). Adapted from Teyssier (2017c).

Continuum absolute calibration accuracy

The continuum calibration is in principle insensitive to sideband gain ratio, so that this uncertainty should not enter the final budget for continuum measurements. The same is true for contributions from the residual optical standing waves when they are corrected through further post-processing. However, the continuum is very sensitive to instrument response drift, which in the end is the main contributor to the continuum flux uncertainty. This effect is discussed in Section 5.8.3.

On top of that, we recall that the intensity calibration of Level 2 and Level 2.5 products for Pointed and Mapping modes is optimised for lines, but not for continuum. Because of the DSB nature of HIFI, the continuum in those products is approximatively twice the real one. The exact correction factor depends on the sideband gain ratio applying to the frequency of interest. We provide in Section 5.8.5 a recipe to recover the exact continuum intensity from a DSB spectrum provided by HIFI. Note that this calibration aspect does not apply to Spectral Scan data since they are already calibrated back on an SSB scale where both line and continuum intensities are treated correctly.

Finally, users should remember that Frequency Switching observation cannot recover the continuum so that this information in not reliable in this type of data.

What if I have removed the standing waves in my spectra ?

Beware that removing standing waves based on a fit of their modulation of the continuum baseline does not necessarily mean that they are corrected in the spectral channels where lines are featured. Indeed it is assumed that the residual standing wave amplitude is proportional to the input signal so that the subtracted amplitude fitted on the continuum modulation does not necessarily match that applying to the line profile. In that sense, the uncertainty component assigned to residual optical standing waves should still be considered for strong lines even if apparent baseline modulations have been removed.

In-band calibration accuracy

The in-band calibration accuracy is driven by any IF-dependent effects that would not be taken into account in the assumed calibration parameters. The main component of this is the sideband gain ratio, for which we only provide IF-dependent calibration in the first $\sim$ 40 GHz of band 2a (Teyssier & Higgins 2013, Kester et al. 2016). In all other bands, the sideband gain ratio is assumed to be flat on the IF. We estimate that additional in-band uncertainties of at most $\sim$ 5% can apply in band 1 from one end of the IF to the other, i.e. the uncertainty is less for data at channels being less separated on the IF. In all other bands, this effect is below 2%.

Calibration accuracy and pointing

The numbers presented in this section apply to a perfectly pointed observation, i.e. a source perfectly centred on the commanded telescope aperture. In practice, the pointing will always suffer from a combination of random and systematic errors that will lead to certain flux loss. This is described and quantified in more detail in Section 5.6.

5.8.3 Relative calibration accuracy

We call relative calibration accuracy the reproducibility figure of a given line or continuum intensity over repeated observations of the same line-of-sight at the same frequency. This figure has been assessed using various datasets obtained repeatedly over the mission and that are described in Teyssier (2017c). The following subsections summarise the main results.

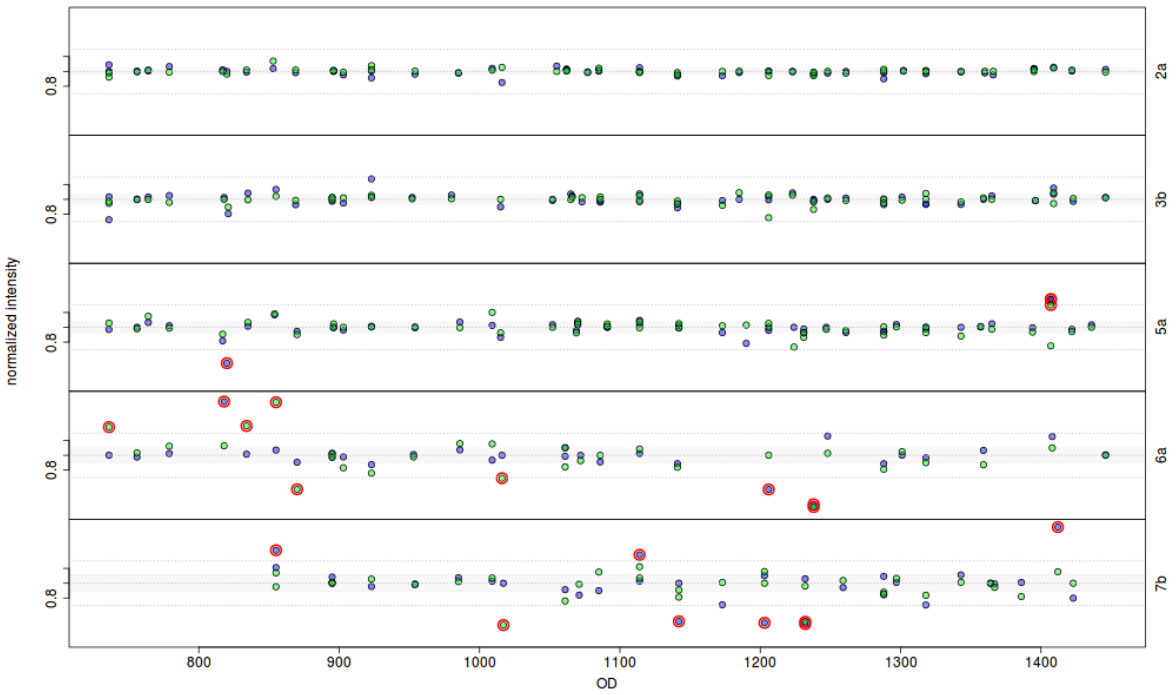

Figure 5.41: Reproducibility of spatially fitted intensities per band and source. Intensities were normalised by dividing them by the median for a given combination of source, band and polarization (from Teyssier 2017c).

Line intensity repeatability

The line intensity relative calibration was obtained from a set of repeated observations towards the circum-stellar envelope of Asymptotic Giant Branch (AGB) stars (Olberg 2015). Within the HIFI beam, these targets are point-like or semi-compact and they were observed through small 3×3 Nyquist sampled raster maps.

Figure 5.41 illustrates the data collected at various epochs towards the set of AGB stars. A range of line intensity ranging from 10 to 300 K km/s is probed by this dataset. It is possible to demonstrate that the line intensity scatter with time is proportional to the source line strength, which allows to treat the repeatability as a multiplicative factor that we provide as a percentage in Table 5.7. The repeatability figures are very good, and range from 3 to 12%, with poorer performances in the HEB bands, typically because of their poorer stability. Note that these numbers not only reflect the overall reproducibility of the intensity calibration, but also the effect of random pointing errors within the *Absolute Pointing Error* (see Sánchez-Portal et al. 2014). It is possible to correct from this latter using the spatial information offered by the 3×3 rasters and reconstruct the intensity at the peak of the map. When those numbers are used to check the repeatability, an improvement of the relative calibration uncertainty is observed, as shown in the last row of Table 5.7. The difference with the non-corrected numbers is a representative measure of how much random pointing errors contribute to the reproducibility of HIFI line data, at least for this type of source morphology.

Source	Band 2a		Band 3b		Band 5a		Band 6a		Band 7b	
	H	V	H	V	H	V	H	V	H	V
CRL 2688	2.2	0.9	3.3	1.5	3.4	3.9	0.9	5.3	10.3	3.8
CRL 618	3.9	2.9	11.2	8	13.7	15.9	NA	NA	NA	NA
IRC 10216	1.1	0.9	1.7	2.4	3.6	2.8	6.6	11.8	10	9.7
NGC 7027	5.3	1.5	7.1	2.6	5.2	4.2	6.3	14.5	NA	NA
o Ceti	3.7	1.4	4.3	2.1	9.7	12.5	9.3	16.6	1.9	13.3
R Dor	3.7	2.4	5.6	4.2	5.1	7.2	8.3	21.1	14.3	12.5
VY CMa	2.1	2.5	4.3	2.1	3.3	5.7	14.4	1.5	NA	NA
all sources	4.3	3.3	8.3	5.6	6.5	8.8	9.3	11.7	12.1	11.2
all sources (ptg-corr)	2.7	2.6	3.5	3.1	6.3	8.1	9.6	9.8	10.1	8.3

Table 5.7: Relative line intensity calibration accuracy derived from the AGB dataset, in percentage of normalized intensities per band and source. NA means that no data were available in this band and source. The global relative uncertainty is provided both for the sample of intensities observed at the commanded source position, and for the sample of intensities corrected for the potential mis-pointing applying to each observation (see Teyssier 2017c for further details).

Continuum intensity repeatability

While the dataset collected on AGB stars was primarily dedicated to the line intensity repeatability study, no equivalent dedicated observing program was foreseen to assess the repeatability of the continuum. It is possible to extract the continuum from some of the AGB star spectra and verify its reproducibility, however the fluxes are fairly weak (2 to 30 Jy at 500 µm, 5 to 130 Jy at 250 µm), especially in comparison with the time spent on each point, making the signal-to-noise ratio of these continua significantly poor. Weak continua are, however, detectable, and have good repeatability when the measurement error is not dominated by the noise (see example of Ceres continuum repeatability thereafter, and Müller et al. 2014). There are, however, other datasets well suited to study the reliability of the continuum calibration over repeated measurements, namely:

- Repeated observations on solar system objects, in particular on Mars and Ceres. Measurements towards Uranus, Neptune and Jupiter are also used to compare to the expected model fluxes. These data proceed from regular beam calibration measurements (see Section 5.9) and from the HSSO Guaranteed Time Key Programme (e.g. Hartogh et al. 2009).
- Repeated observations from science programmes focussing on absorption studies towards the diffuse ISM – most of those have been collected in the framework of the PRISMAS Guaranteed Time Key Programme (e.g. Gerin et al. 2010, De Luca et al. 2012) and follow-up studies.

We provide in the following the outcome of those studies (This section is still to be written).

5.8.4 Overall H and V flux intensity imbalance

As described in Appendix A and Figure 5.44, the respective H and V polarisation apertures do not strictly overlap on the sky and may therefore pick up slightly different flux from the pointed area of

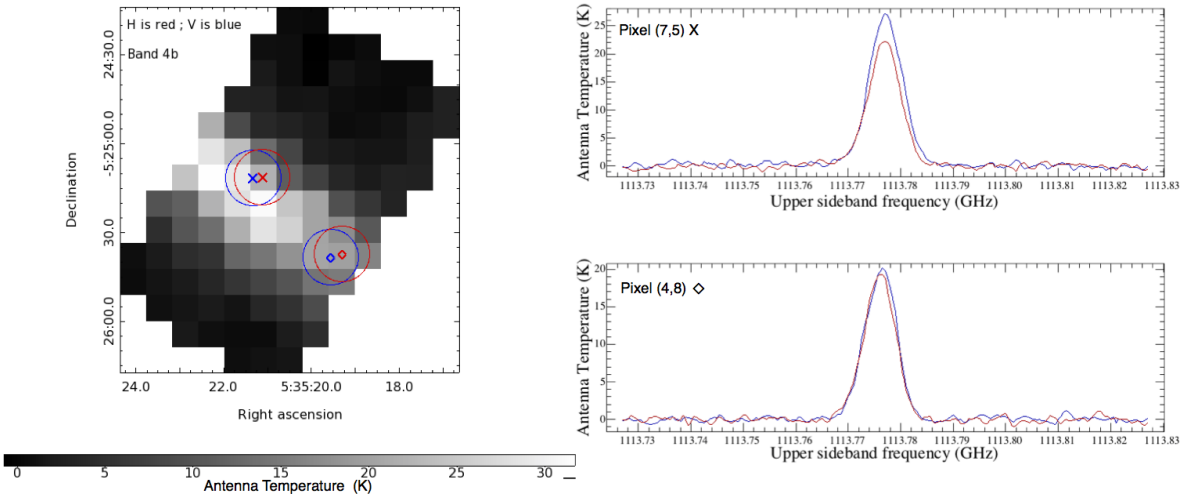

Figure 5.42: Illustration of severe H/V imbalance in the presence of steep gradients in the source brightness distribution. *Left*: 1113 GHz water line map (cube layer at the line peak intensity channel) of the Orion Bar. Red and Blue symbols show the respective positions of the H and V apertures (with their HPBW) at two different locations on the Bar. *Right*: corresponding H (red) and V (blue) spectra at the two locations shown on the map. Note the strong H/V imbalance in the Northernmost position, while the two polarisations agree remarkably well at a location showing a flatter source morphology more at the South.

the sky. This will lead to unavoidable intensity discrepancies between the spectra obtained in either polarisation, although these differences should be small – Section 5.6 gives an overview of the typical impact of the pointing on the H/V line and continuum intensity ratio. There are, however, other instrumental effects that can lead to imbalances between the two polarisations, namely:

- differences in the sideband gain ratio: these should be corrected for in the pipeline
- differences in the beam coupling: these are not corrected for in the data calibrated in T_A^* and they can account for between 3 and 5% depending on the band (see Table 5.8).

Overall, the bulk of the H/V intensity imbalance can be explained by pointing offsets (Section 5.6). The most drastic impact of this effect will be observed for exceptional cases where the source brightness distribution exhibits steep gradients on very small spatial scale. Such cases have been observed for example in maps of the Orion Bar, as illustrated in Figure 5.42.

5.8.5 Continuum intensity in double sideband data

We provide here a simple recipe in order to recover the applicable continuum intensity in a DSB spectrum provided at Level 2 or Level 2.5 in the Pointed and Mapping observing modes (this is not needed for Spectral Scan observations). The bottom line is that the sideband ratio correction that is needed for the accurate line calibration has to be undone. If one calls $C_{\rm DSB}$ the continuum level

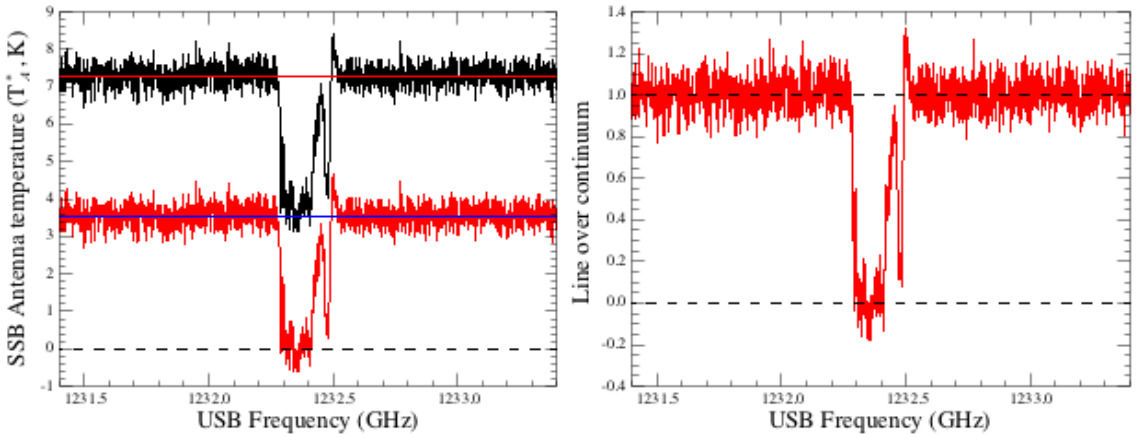

Figure 5.43: Illustration of the correction of the continuum level in a spectrum calibrated to a SSB scale by the pipeline. The example shows the HF line seen in absorption in W31C (Neufeld et al. 2010) at the upper end of band 5, where a $G_{\rm ssb}$ value of 0.485 applies for lines in the USB. *Left*: Level 2.5 spectrum as provided by the HIFI pipeline (truncated to the 2 GHz around the HF line) in black. In this spectrum, the measured continuum $C_{\rm DSB}$ (shown as a red horizontal line) corresponds to ~ 2.06 times the real SSB continuum. The SSB spectrum after the continuum correction following Eq. 5.3 is shown in red, with its corresponding continuum drawn as the blue horizontal line – note that the line intensity is not changed by this transformation, simply its contrast with respect to the effective SSB continuum. *Right*: corresponding line-over-continuum spectrum after continuum correction, showing that the line is very likely saturated.

estimated in the Level 2.5 products of either a pointed or mapping mode observation, and $G_{\rm ssb}$ the normalised sideband gain ratio applied to this dataset[2], the SSB continuum value is simply given by:

$$C_{\rm SSB} = C_{\rm DSB} \times G_{\rm ssb} \tag{5.2}$$

This simple correction assumes that the continuum changes slowly with frequency, i.e. it takes very similar values in the USB and LSB ranges respectively. Note that this is not strictly true in case of a steep dust spectrum but this is a reasonable assumption is the vast majority of the cases.

When both line and continuum need to be well calibrated on the same spectrum (e.g. to assess the depth of absorption lines), the following transformation is needed on the Level 2 or Level 2.5 spectrum $T_{\rm SSB}(\nu)$:

$$T_{\rm SSB}^{\rm corr}(\nu) = T_{\rm SSB}(\nu) - [C_{\rm DSB} - C_{\rm SSB}] = T_{\rm SSB}(\nu) - C_{\rm DSB}(1 - G_{\rm ssb}) \tag{5.3}$$

This correction is illustrated on 5.43 for the HF line seen in absorption at the upper end of band 5 – see caption for further details.

[2] the value of this parameter can be found in the product header under the keyword `USBGAIN` or `LSBGAIN` depending on the applicable sideband

5.8.6 Conversion to other scales

Users of the HIFI Level 2.5 products will normally need to convert the data calibrated onto the T_A^* scale into a more physical and instrument-independent scale. The decision of which transformation to apply depends both on the source brightness distribution size with respect to the HIFI main beam, and the kind of science analysis approach that is foreseen with the data.

If users plan to perform their scientific analysis remaining on a Rayleigh-Jeans equivalent temperature scale, the transformation from the T_A^* scale is mainly aimed at correcting for the flux loss associated to the imperfect coupling of the telescope beam with the source brightness distribution. It is a common practice among radio-astronomers to convert T_A^*-calibrated data into the main beam temperature $T_{\rm mb}$, following the simple formula:

$$T_{\rm mb} = \frac{\eta_l}{\eta_{mb}} T_A^* \tag{5.4}$$

where η_l is the forward beam efficiency (assumed to be 0.96 for all HIFI bands and frequencies), and η_{mb} is the main beam efficiency, whose values are summarised in Table 5.8, and described in more detail in Mueller et al. (2014). Strictly speaking this transformation is only recovering the exact flux when the source uniformally fills the instrument main beam (Section 4.4.2), however, it is considered a good approximation for source sizes "close" to that of the instrument main beam. When the source brightness distribution is larger than the beam, the real brightness temperature will be somewhere between T_A^* and $T_{\rm mb}$, where T_A^* would only be the correct scale of a source filling the whole forward hemisphere. Although users often assess their source sizes from e.g. integrated intensity maps, one should be reminded that the source brightness distribution per velocity channel usually fills a lesser fraction of the sky than that of the integrated intensity maps. Since data are calibrated per spectral resolution channel, the $T_{\rm mb}$ remains a good approximation even for slightly extended source morphology. Dealing with cases that would not fulfil this approximation involves ad hoc source brightness distribution models that need to be convolved with the exact, applicable HIFI beam (Mueller et al. 2014) and is beyond the scope of this document.

When users typically need to compare their HIFI data to data obtained from other instruments (radio-telescopes or not), the best practice consists in converting the brightness temperature scale into a flux density scale in Jansky. The conversion from Kelvin to Jy is already described in Sections 4.4 and 4.4.1, and is given by the general formula (rewriting Eq. 4.6):

$$S_{\nu,tot}(Jy) = \frac{\eta_l}{\eta_A}\frac{2k}{A_g}\frac{\Omega_S}{\Omega_\Sigma} T_A^* = \frac{\eta_l}{\eta_A}\frac{2k}{A_g}\frac{1}{K} T_A^* \tag{5.5}$$

where η_A is the aperture efficiency, whose values are summarised in Table 5.8, and described in more details in Mueller et al. (2014), A_g the telescope geometric area, k the Boltzman constant, and $K = \Omega_\Sigma/\Omega_S$ the flux dilution factor due to beam effects (Ω_S is the source solid angle, and Ω_Σ the beam-weighted source solid angle, see also Jellema 2015). In the limit of point sources, K=1, while for non-point sources, $K < 1$. When all constant parameters are merged, the conversion factor simply writes $S_{\nu,tot}/T_A^* = 313/(\eta_A K)$, i.e. for a point source, there are typically 470–560 Jy per measured Kelvin. For non-point-like sources, users will have to make an assumption about the typical source size.

There exist analytical approximations of this conversion factor in the case of purely Gaussian beams and for either Gaussian or top hat (uniform circular sources) source brightness distributions. They are discussed in Kramer (2008). In the case of HIFI, we show in the next section that the beam cannot

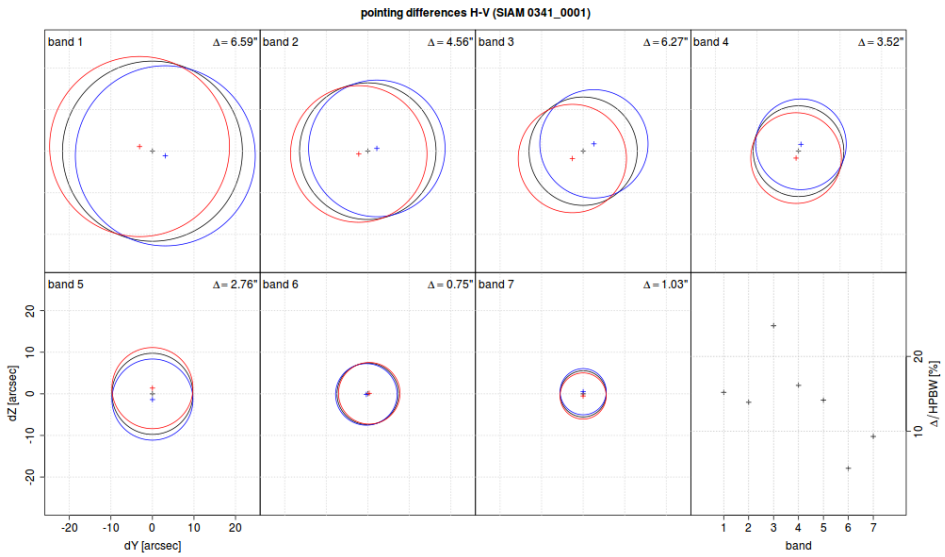

Figure 5.44: The observed offsets between the H and V polarisations per band. Positions are shown relative to the synthesized beam (black crosses), assumed to be at the mean position of H (blue crosses) and V (red crosses). The distance between H and V in arc seconds is given as Δ. Circles indicate the half-power beam widths (HPBW). The last panel shows the distance between H and V as a percentage of the HPBW for each band.

be assimilated to a simple 2D-Gaussian profile. For users desiring to compute this conversion factor taking into account the exact HIFI beam profile, a dedicated task, `convertK2Jy`, is available in HIPE.

5.9 The HIFI beams and sky coupling efficiencies

5.9.1 HIFI SIAM

Pointing offsets for the various HIFI apertures were determined during the commissioning phase via observations of Saturn. Individual phases of DBS measurements were analysed in order to determine the polarisation and chopper dependence of the SIAM entries for the HIFI instrument (see Appendix A). The inferred H and V respective aperture positions in the HIFI focal plane are shown in Figure 5.44, indicating in particular the H/V co-alignement offset (labelled as Δ) - see also Sections 5.6.2 and 5.8.4 for a discussion on the impact this offset has on line fluxes.

From the fair number of 3×3 DBS maps of various AGB stars in selected HIFI bands (2a, 3b, 5a, 6a and 7b) which were obtained during a routine phase calibration program, pointing offsets could be determined for H and V polarisation separately by fitting 2D-Gaussian distributions to the observed integrated intensities. As an example we show results for the source R Dor in Figure 5.45. When transformed to the focal plane (FP column), the orientation of the H/V separation lines up with what is expected from the SIAM[3]. On ground, extra care was taken to properly align H and V for the HEB

[3]Note, that there is a change of sign involved when going from the SIAM to the focal plane. When moving the telescope

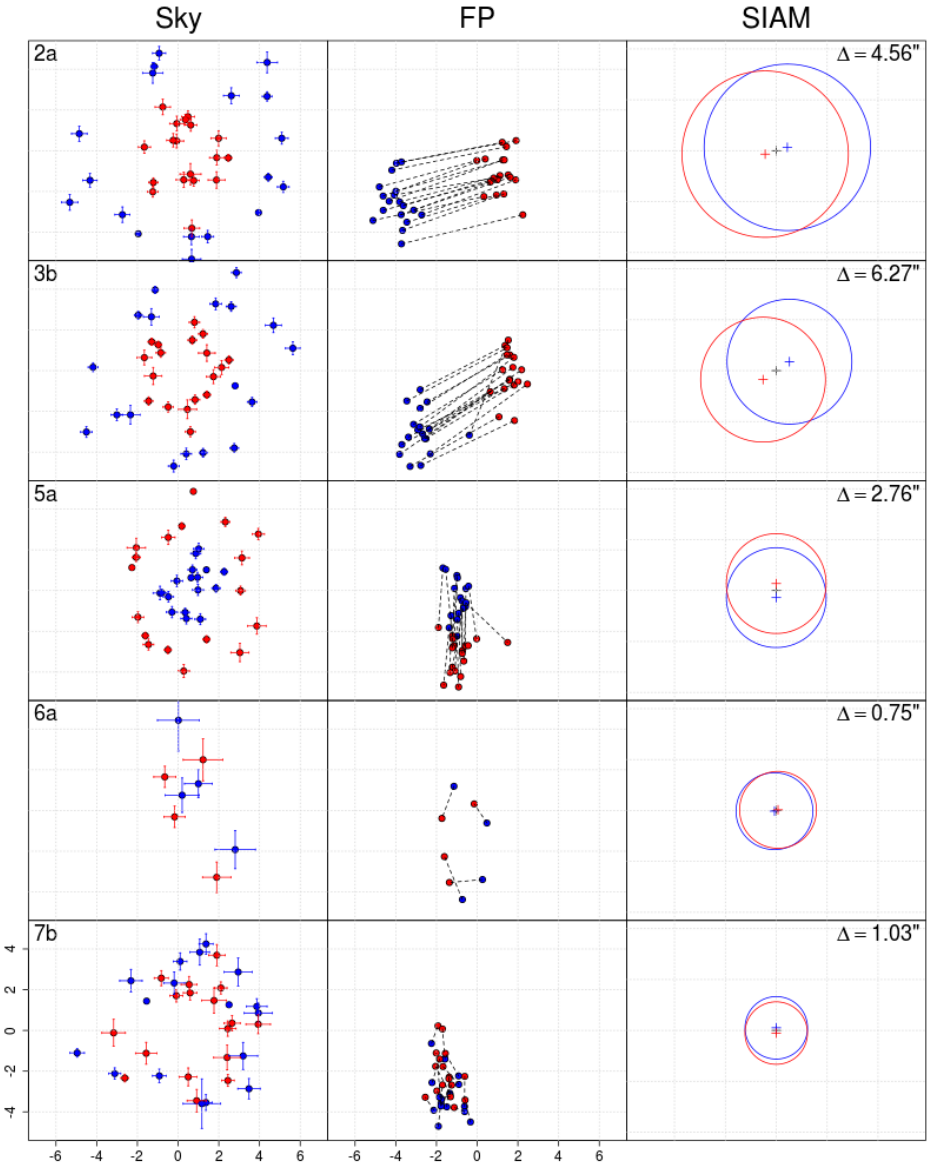

Figure 5.45: Pointing offsets determined from mapping observations of R Dor in various HIFI bands. When transformed to the focal plane (FP) these observations confirm the expected relative orientation of H and V from the SIAM. Colour coding is again blue for H and red for V polarisation.

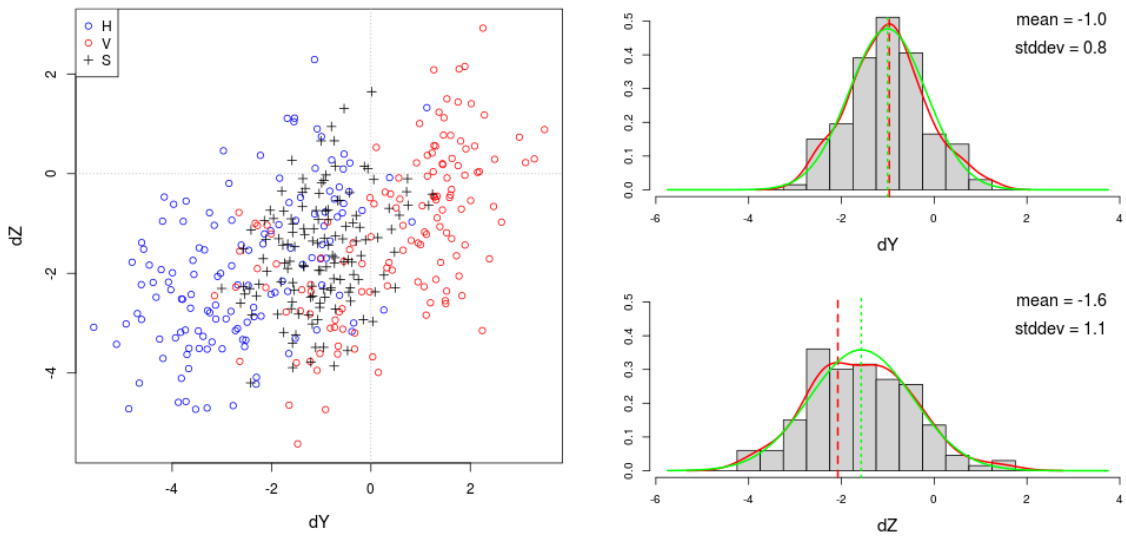

Figure 5.46: Distribution of observed pointing offsets in the focal plane, shown as both a scatter diagram (left panel) and as histograms (right panels).

bands, which are the most sensitive to pointing errors due to their narrow beams. Our routine phase observations confirm that these are well aligned.

Finally, Figure 5.46 shows the distribution of the observed pointing offsets. There appears to be a remaining bias of $-1''$ and $-1.6''$ along the Y and Z axes of the focal plane coordinate system, respectively.

5.9.2 HIFI beams

To characterise the HIFI beam, the ICC performed in-flight mapping observations of Mars, a bright semi-extended source of known brightness and size, using all HIFI bands. Two "runs" of such Mars observations were performed in 2010, soon after launch. Roelfsema et al. (2012) derived HIFI's half-power beam width (HPBW), aperture efficiency η_A , and main-beam efficiency η_{mb} from these data, assuming a Gaussian beam shape, as is customary for radio observatories. Two additional "runs" of Mars observations were performed later into the mission.

A Gaussian beam model provided an unsatisfactory fit to the data. Instead, we derived detailed 2D beam models by accounting for obscuration and truncation by the secondary mirror M2, the obscuration by the tripod supporting the secondary, and using the measured trefoil-shaped wavefront error of the primary mirror M1. The resulting beam models show significant side lobes. At high frequencies, azimuthal symmetry is visibly broken: the tripod supporting M2 induces an approximately triangular beam shape. At lower frequencies, beams are "rounder". Notably, the new beam models tend to be narrower than the previous Gaussian estimates, typically by a few percent.

to the right, objects in the field of view move to the left.

Mixer	f GHz	η_{mb}	η_A	HPBW $''$	K-to-Jy Jy/K
1H	480	0.62	0.65	43.1	482
1V	480	0.62	0.63	43.5	497
2H	640	0.64	0.64	32.9	489
2V	640	0.66	0.66	32.8	474
3H	800	0.62	0.63	26.3	497
3V	800	0.63	0.66	25.8	474
4H	960	0.63	0.64	21.9	489
4V	960	0.64	0.65	21.7	482
5H	1120	0.59	0.54	19.6	580
5V	1120	0.59	0.55	19.4	569
6H	1410	0.58	0.59	14.9	531
6V	1410	0.58	0.60	14.7	522
7H	1910	0.57	0.56	11.1	559
7V	1910	0.60	0.59	11.1	531

Table 5.8: Adopted values for η_{mb}, η_A, HPBW, and point-source sensitivity S/T_A^* (i.e. Kelvin to Jansky conversion factor) $= (2k_B/A_g)(\eta_l/\eta_A)$ for one spot frequency per mixer. This Table is to be compared to Table 5 in Roelfsema et al. (2012), which uses the same spot frequencies.

The data and results from the beam modelling were released in October of 2014 (Mueller et al. 2014). The delivery consisted of 28 FITS files, together with this note. The exact derivation of the beam models are described in Jellema (2015).

We provide in Table 5.8 the adopted values of η_A, η_{mb}, HPBW, and factors of $2k_B/(\eta_A A_g)$ to convert from Kelvins to Janskys (see (4.6) in Section 4.3) for one spot frequency per mixer; we chose the same spot frequencies as in Roelfsema et al. (2012, Table 5).

As part of the beam release, HIFI provided beam models for two spot frequencies per HIFI mixer band, separately for polarisations H and V, i.e., a total of $7 \times 2 \times 2 = 28$ beam models. Spot frequencies are chosen to be close to the central frequency of LO subbands a and b, respectively.[4]

Each model was released as a FITS file, each containing three extensions:

- "Primary" containing a 2705×2705-pixel 2D beam model
- "Azimuthally averaged beam profile" containing a table of the azimuthally averaged beam as a function of radial distance to the centre (spatially oversampled compared to the beam itself)
- "Encircled Energy Fraction (EEF)", the integral of the above.

As an example, two 2D beam models are plotted in Fig. 5.47. A standard world coordinate system (WCS) is given in the FITS header of the primary extension. It is given in focal-plane coordinates (see Sect. 5.9.3); axis 1 corresponds to angular offsets along the spacecraft Y axis, axis 2 to angular offsets along the spacecraft Z axis. The beam intensity is given in linear units, peaking at the centre of the

[4]We emphasise that, for our current purposes, the LO subbands are merely used to label frequency ranges. There is no reason why the local oscillator itself should have any influence on the beam shape, which is a function of the telescope and the mixer.

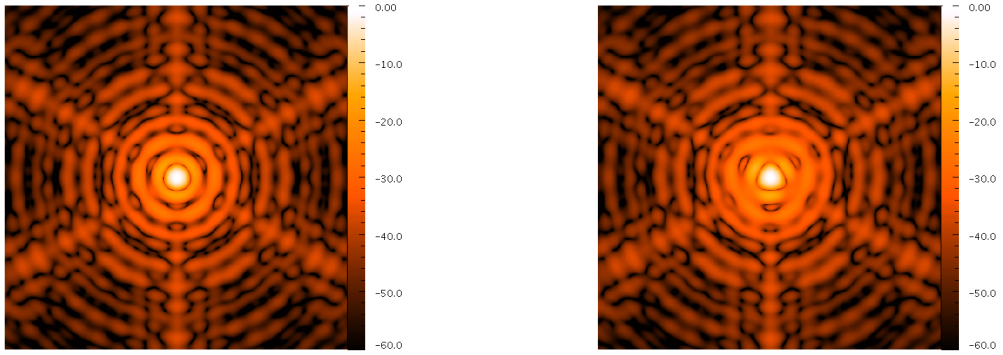

Figure 5.47: Two sample beam shapes: band 1 in polarisation H at 520.5 GHz; band 7 in polarisation V at 1847.5 GHz The beam intensity is colour-coded (in dB), normalised to unity (0 dB) at the centre. All beam models display appreciable azimuthal asymmetry, more so in the outer regions. The central main beam is reasonably round at low frequencies, but shows increasing asymmetry with increasing frequency. Here we show the central 541 pixels out of 2705, i.e., the complete beam models are larger than what is depicted here.

image (`CRPIX1` and `CRPIX2`, note that both are one-based indices following the FITS convention, not zero-based) and normalised to unity at the peak.

The spatial resolution of our beam models scatters around 18 points per HPBW. The provided FITS files therefore cover a square spatial range from roughly -75 to $+75$ HPBW. This constitutes a significant extrapolation beyond the range covered by our mapping observations (which typically went out to ~ 2 HPBW).

This is motivated in part by the slow convergence of the encircled-energy fraction (see discussion below): even rather large spatial distances add significant power to the HIFI beam. Note, however, that typical HIFI observations contain an `OFF` or a `REF` position, assumed to be blank sky, which is subtracted from the on-target data. This mitigates effectively against contamination from large-scale structure that may be picked up by the error beam: **only the flux gradient over length scales determined by the on-off distance remains as contaminant**.

Full-size beam models are provided so users can extract the best possible science from their data. For most applications, however, it is enough to crop the beam models as appropriate, significantly reducing the computational cost of the required numeric convolutions. To scale the beam model to a frequency f, the pixel scale given in the FITS header (`CDELT1` and `CDELT2`) should be multiplied by f_0/f, where f_0 is the spot frequency for which the delivered beam model is valid.

5.9.3 Rotation to sky coordinates

HIFI beams are naturally expressed in a spacecraft-fixed coordinate system, seeing that their spatial structure is caused by various components of the optical path. We use the conventional orthonormal Herschel-fixed coordinate system (see the Herschel Observers Manual, and references therein) in which the positive X axis is aligned with telescope boresight, the positive Z axis points to the centre (in width) of the sun shield, and the Y axis follows from the right-hand rule.

The spatial coordinates of the HIFI beams are angular offsets along the positive Y and Z axes, off the nominal position of the active mixer. These coordinates are called "focal-plane coordinates" in the following. The conversion between angular offsets in the focal plane $\mathrm{d}Y, \mathrm{d}Z$ and along the RA and Dec axes $\mathrm{dRA}, \mathrm{dDec}$ depends on the telescope roll angle, β (column `posAngle` in HIFI datasets):

$$\begin{pmatrix} \mathrm{dRA} \\ \mathrm{dDec} \end{pmatrix} = \begin{pmatrix} \cos\beta & \sin\beta \\ -\sin\beta & \cos\beta \end{pmatrix} \begin{pmatrix} \mathrm{d}Y \\ \mathrm{d}Z \end{pmatrix}. \tag{5.6}$$

Telescope boresight corresponds to $\mathrm{dRA} = \mathrm{dDec} = 0$. To linear order, RA coordinate offsets need to be multiplied by $\cos$ (Dec) to get angular offsets.

Chapter 6

HIFI observations in the Herschel Science Archive

In this chapter we briefly describe the data products and the data structure coming from HIFI observations. We provide an overview of the science readiness of the HIFI data, as well as recommendations and recipes in case any further data processing still applies. A comprehensive explanation, and practical guidelines for accessing the data, running the pipeline and doing quick interactive analysis can be found in the HIFI Data Reduction Guide, which can be accessed also from the help documentation inside the Herschel Interactive Processing Environment (HIPE).

The data structure for the *Herschel* mission may be unusual for astronomers as it is object-oriented – the underlying pipeline processing and the database that hold all the data are based on Java. No other ground-based or space observing facility provides similar access to the collected data up to now. In some situations the data access and the data structure may seem confusing. In such cases please contact the helpdesk of the Herschel Science Centre.

6.1 HIFI processing levels and structure of the observation context

All standard *Herschel* HIFI observations are now publicly available in the Herschel Science Archive (HSA). Getting observations from the HSA is explained in greater detail in Section 1.4 of the HIPE Data Analysis Guide.

All *Herschel* HIFI data come from the HSA a directory-tree structure organised by procesing level. When inspected in HIPE, those levels are referred to as context. A context is a special kind of product linking other products in a coherent description and can be thought of as an inventory or catalogue of products. The HIFI processed observation consists of many such contexts enclosed within one *Observation Context*. The Observation Context is therefore a product container, which is comprised of various layers of science data, auxiliary data, calibration data, and other information about the observation in an onion skin type of structure.

Each HIFI Observation Context is the outcome of the Standard Product Generation (SPG) run in bulk with a given version of the *Herschel* Common Software System (HCSS). The SPG version, present in the metadata of any Observational Context, reflects the HCSS version of the reprocessing. The HIFI Science data are found in the Level 0, 1, 2, and 2.5 Contexts and are the result of each stage of the

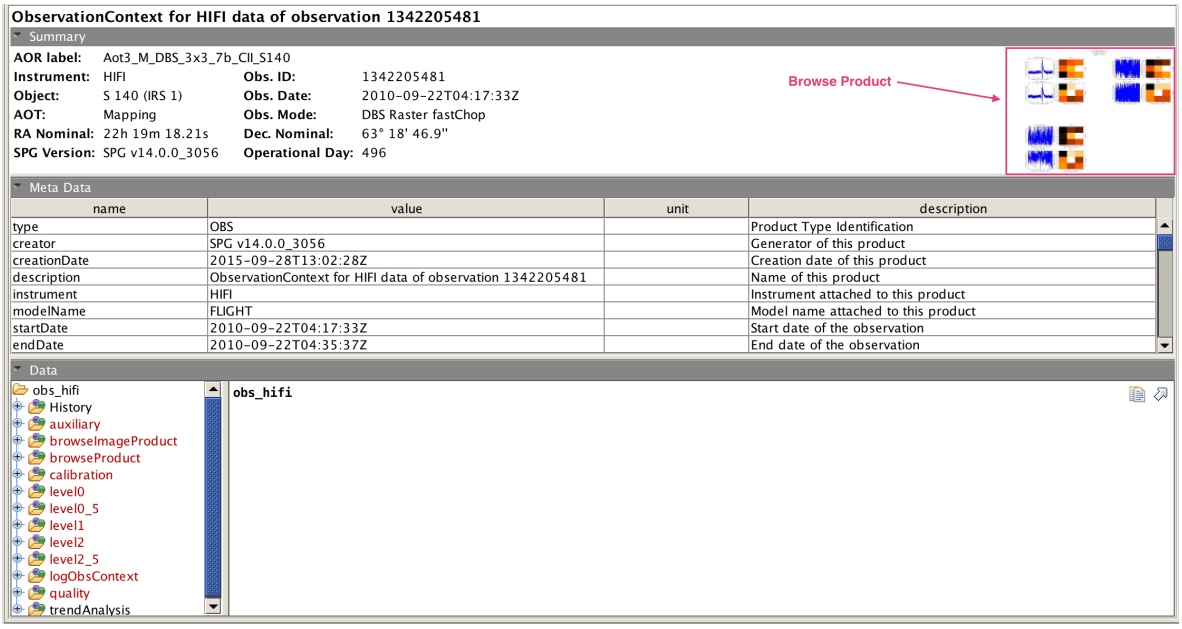

Figure 6.1: An example of the ObservationContext structure of a HIFI observation.

pipeline (data at Level 0.5 are removed in order to conserve memory as they are less relevant to the archive user). For HIFI, the observation Context contains the following: (see also Figure 6.1)

- **Level 0:** The Level 0 Context contains a *HifiTimelineProduct* (HTP) for each spectrometer used in the observation, which in turn contains a dataset for each integration in the observation, and as a consequence contains more datasets than higher levels of processing. The Level 0 Context also contains a Quality Product based on checks on the Level 0 data for telemetry issues common to both the WBS and HRS (in the *CommonTm* product), and for checks on the data frame count and quality (in the Quality product for each spectrometer).

  Level 0 is the most raw that you will see your data. It has been minimally manipulated into a HTP, had pointing information from the satellite associated with it, and undergone several ”sanity checks” to flag any incidences of housekeeping parameters. Additionally, information from the *Uplink* product, which contains information calculated by HSpot concerning how the observation should have been carried out (e.g., the spacing between scan legs in a map) is copied to the HTP and metadata. Level 0 data has units of channel number (”wave” scale) and counts (intensity scale). As such those products are of very unlikely interest to most of the archive users.

- **Level 0.5:** This stage of the pipeline focusses on removing the spectrometer instrumental effects – essentially this is where the frequency calibration is performed. There are separate pipelines for the WBS and HRS spectrometers. This product is actually absent from the Observation Context as soon as the Level 1 products are successfully generated. It can, however, be re-generated if needed (see 6.4.1). Some Quality information is retained at Level 0.5, and contains the WBS-H and WBS-V comb, and zero quality checks.

- **Level 1:** The Level 1 context contains HTPs for each spectrometer used in the observation and a Quality product. Level 1 data are frequency and intensity calibrated, and are also corrected for the velocity of the spacecraft. The application of observing mode specific calibrations, i.e., subtraction of reference and OFF positions, and intensity calibration using Hot/Cold loads, is done at that stage. Data processed to Level 1 are in the frequency IF scale (in MHz) and on the T'_A temperature scale (in K), i.e. it does not yet take into account losses due to the rearward beam (see Section 4.3).

  The quality product at Level 1 contains, for each spectrometer used in the observation, the results of the phase checks done by the Level 1 pipeline.

- **Level 2:** The Level 2 Context contains an HTP for each spectrometer used in the observation. Level 2 data are converted to antenna temperature scale T^*_A (in K) and to sky frequency (GHz). Due to the DSB frequency degeneracy, products are generated both on a USB and on an LSB frequency scale. Spectra are averaged together, per each spectrometer, for each LO setting, and each spatial position in the observation. This results in a single spectrum (for each spectrometer) for point observation mode, individual LO setting for spectral scans, and individual spectra per position and LO setting for maps. These data products may be at a publishable quality level, although corrections for baseline issues are likely required. Electric Standing Wave have been corrected at this point for all modes.

- **Level 2.5:** The Level 2.5 pipeline combines the Level 2 products into final products and depend on the observing mode:

  Point mode spectra are stitched, folded (in the case of Frequency Switch), and converted to SimpleSpectrum format. HRS data are stitched only in the case that subbands overlap in frequency.

  Mapping mode Level 2.5 HTPs contain spectra that are stitched and folded (in the case of Frequency Switch), a cubesContext is also produced containing cubes constructed from the stitched and folded (if applicable) spectra. Stitching of HRS data is done as for *Point mode* observations. In the case of mapping observations carried out with a non-zero rotation angle, a cubesContextRotated is also produced, which contains cubes generated from the Level 2.5 HTPs and with the rotation angle applied.

  Spectral Scan Level 2.5 contains a deconvolved single sideband spectra per polarisation.

- **Calibration Context:** The CalibrationContext contains the data passed to the pipeline for calibration (*Downlink*), the calibration files created by the pipeline (*Pipeline-out*), and information about how the observation was carried out (*Uplink*). Section 6.5 provides further details on this context.

- **Auxiliary Context:** This context contains mostly products that were not generated by HIFI (e.g. Orbit files, Ephemerides, Spacecraft-related products such as the pointing system – the ACMS, etc), however, this is also where the *HifiUplinkProduct* is stored. This product contains information about how the observation was carried out, reflecting the outcome of the sequencer computation in HSpot (Section 3.2).

- **Quality Context:** This Context contains information about the Quality Flags that were raised automatically by the pipeline, as well as the processing log. It is extended in an additional context called the *Quality Summary Context* which is generated manually after inspection of

the data, and contains complementary information on the data quality. More details can be found in Section 6.3.2.

- **Browse Product:** This product offers a quick look at the data through automated extraction from the results of the *standard* pipeline. It corresponds to the postcard displayed in the HUI and shows the outcome of the Level 2.5 product generation. Note that those products are not necessarily science-ready.
- **Trend Analysis Context:** This Context contains products useful for tracking systematic changes in instrument response over time. Of possible interest to the archive user are the tables stored in the Statistics tables, which contain information about the spectra moments such as mean and noise RMS among others.

If you are interested in an in-depth description of the Observation Context, the main reference is the HIFI Data Reduction Guide. For further details of individual pipeline tasks, the reference is the HIFI Pipeline Specification Document. Finally, the detailed steps and algorithm implemented in the HIFI pipeline are described in Shipman (2017).

6.2 HIFI products in the HSA

As global repository for the Herschel mission, the *Herschel* Science Archive contains all the information collected and processed during the mission. On top of the SPG output, it also contains quality-enhanced and added-value products, as well as a certain number of Ancillary Data Products, not necessarily associated to observations, which we also describe thereafter.

6.2.1 Standard Pipeline Products

The products generated by the instrument pipeline are described in the previous Section 6.1. They correspond to the latest bulk reprocessing version (which for HIFI is that of HIPE 14.1), although older product versions can also be fetched in the HSA. See the previous section and Shipman (2017) for more details.

6.2.2 User Provided Data Products

User Provided Data Products (UPDPs) are sets of products delivered by observing time holders. In the vast majority those will correspond to data collected in the framework of Guaranteed Time programmes. Since some of those products have been delivered already quite some time ago, they do not always reflect the latest knowledge about the pipeline processing or data calibration. As such they should be used with caution and it is strongly recommended to carefully read the release notes associated to each of those deliveries.

The list of available products is given in *Herschel* UPDP page. Note that UPDPs are not necessarily limited to spectra data, they can also contain catalogues, models, or combinations of data products not contemplated by the standard pipeline. Specially curated UPDPs can also be provided by the Herschel Science Centre. They are summarised in Table reftbl:hsc-updp. Presently the only dataset in that category is that of the non-averaged Level 2 spectra for pointed mode observations. These latter

Product name	Description	Deliverables	Release note
Non-averaged Level 2 data products	Individual spectra from the Level 2 processing prior to the average performed otherwise to form the Level 2.5 products from the pipeline	Level 2 Spectra (HTP)	Teyssier (2016c)

Table 6.1: List of User-Provided Data Products curated by the Herschel Science Centre

HPDP name	Description	Deliverables	Release note
Spectral Scans	baseline-cleaned deconvolved data	Deconvolved spectra, postcards	Teyssier (2017a)
Spectral Maps	baseline-cleaned cubes, optimally regridded. These products also offer integrated intensity maps around the strongest line of interest in a given observation	Cubes and integrated intensity maps, postcards	Morris (2016)
OFF positions	these are post-processed versions of the pipeline-generated OFF position spectra. In particular, spectra are stitched. For spectral scans, deconvolved data are provided, and for DSB raster maps, OFF positions cubes are prepared	Spectra and cubes, postcards	Teyssier (2016d)
Uncertainty tables	Extracted uncertainty tables from the Calibration Context of the HIFI SPG output	Table Datasets as FITS files	Teyssier (2016f)
Line lists	Line identification catalogues for each observation	Catalogues, postcards	Benamati & Caux (2016)

Table 6.2: List of contemplated HIFI HPDPs

allow users to perform their own average of the Level 2 spectra, and potentially discard individual datasets

6.2.3 Highly Processed Data Products

Highly Processed Data Products are sets of products generated by expert scientists, generally from Herschel Science Centre, the NASA Herschel Science Center, and the Instrument Control Centres.In the case of HIFI, emphasis is given on those datasets where particular post-processing is deemed necessary in order to produce science-ready data, although some added-value products are also provided. The HPDPs contemplated by HIFI are summarised in Table 6.2.

HPDPs can be fetched at the following link http://www.cosmos.esa.int/web/herschel/highly-processed-data-products. In the very near future, they will also be served by a dedicated panel in the web-based interface of the HSA (http://archives.esac.esa.int/hsa/whsa/index.html#home).

Product name	Description	Deliverables	Release note
	Science Data		
Gas cell ILT database	HIFI Gascell measurement collected during the pre-launch campaign	Level 2 and deconvolved spectra	Olberg (2017)
HIFI beams	2-D beam maps, 1-D azimuthal beam averages, and Encircled Energy Fractions (EEF) for each band	FITS files of the aforementioned quantities	Teyssier (2016a) and Mueller et al. (2014)
Post-Helium HIFI WBS SEU checks	A time series of WBS calibration spectra aimed at detected Cosmic Ray hit statistics	Spectra	Teyssier & Osenkopf (2017)
	Model Data		
Mars models	Model of the Mars continuum emission used in beam properties assessment	Mars maps at applicable frequencies and dates	*Herschel* Flux calibrator Models page
	Engineering and House-Keeping Data		
Trend Analysis Data Products	Relevant HIFI HK record with time over whole mission	Ascii files and plots	Teyssier & Edwards (2016)
Diplexer Calibration Products	Engineering telemetry record for measurements aiming at calibrating the HIFI diplexers	Ascii files and plots	Teyssier (2016b)
$T_{\rm sys}$ Spectra	Spectra and plots of System Noise Temperatures from full Spectra Scans	FITS files and plots	Teyssier (2016e)
	Software Information		
OBSW images	HIFI On-Board Software Programme Images	Ascii files	Teyssier (2017b)
LCU Software images	HIFI LCU Software Programme Images	Ascii files	
LOU Safety Tables	HIFI LOU Safety Tables	Ascii files	

Table 6.3: List of contemplated HIFI Ancillary Data Products

6.2.4 Ancillary Data Products

Ancillary Data Products (ADPs) are products not necessarily associated with an observation. Although a large fraction of them cover information of little interest to the archive user, we can highlight the following items of relevance for science exploitation:

- the HIFI beams (see also Section 5.9).
- the Mars models used by HIFI as primary calibrator, in particular to derive its beam coupling efficiencies (Section 5.9 and 4.5)
- the database of HIFI gas cell measurements collected during the Instrument Level Tests. These are laboratory measurements that were taken with HIFI pre-launch. In particular a full methanol survey was performed during this period. Details about this dataset can be found in Higgins (2011) and Olberg (2017)

Table 6.3 lists all Ancillary Data Products associated to the HIFI instrument. ADPs can be fetched at the following link http://www.cosmos.esa.int/web/herschel/ancillary-data-products. In the very near future, they will also be served by a dedicated panel in the web-based interface of the HSA (http://archives.esac.esa.int/hsa/whsa/index.html#home).

6.3 Which Products should I use for Data Analysis?

6.3.1 Science readiness of the HIFI products

In most cases, the HIFI Level 2 and Level 2.5 products offer a scientific quality that is sufficient to directly perform further data analysis (e.g. line or continuum intensity extraction). The one step that will always be necessary in any of the HIFI product, though, consists in passing from the HIFI-centric antenna temperature intensity scale to a scientifically-meaningul scale. This is explained in Sections 5.8.5 and 5.8.6 for both line and continuum science exploitation.

Whenever residual artefacts remain, they are mostly in the form of baseline distortion of various nature (Section 5.3), and for which mitigation recipes are provided in Sections 6.4.1 and 6.4.2. We estimate that about 20% of the pipeline products could still be affected by such baseline distortion. As most of those cases can be easily dealt with using standard baseline correction tools, although those products are not strictly speaking *science-ready*, they can be considered at least *science-friendly*.

Since some of those residual imperfections are actually taken care of in the HIFI HPDPs (Section 6.2.3), the best product to use depends basically on the availability or not of the latter. This will be particularly true for Spectral Scans and Spectral Maps taken in bands 6 and 7, which will account for about half of the residual unruly baselines in the pipeline products. Finally, the HIFI products will usually come with two different spectral resolutions (WBS and HRS data) – we leave it to the user's decision as to which of those data are the most suitable to their science (apart from a couple of exceptions, the data quality will be the same in both spectrometers). The following provides a summary of the products that should be used depending on the observing mode.

Pointed Mode Observation

No dedicated artefact correction has been performed for this mode in the form of HPDPs, and we estimate that about 20% of all the pipeline products taken in this mode could still be affected by residual baseline distortion. One can distinguish the following use cases:

- Unless the data are affected by noticeable platforming or parabolic residual baseline distortion (5.3.1), the Level 2.5 stitched products should be used. Otherwise, Level 2 products are best in order to correct for those residual artefacts on a spectrometer subband basis.

- For users wishing to inspect individual Level 2 products prior to their averaging, the SPG products are not fit, instead the corresponding User Provided Data Products should be used (Section 6.2.4).

Spectral Mapping Observation

While the first use case above can also apply to HIFI spectral mapping data, HPDPs have been generated for a subset of those products, offering a higher quality level than the standard generation products. We recommend the following:

- whenever an HPDP exist for a given Spectral Map observation, this product should be primarily used. It will offer both a baseline-cleaned version of the cubes, with cube grid cell sizes optimised for S/N, and integrated intensity maps centred on cherry-picked species of interest belonging to the observation. HPDPs for Spectral Maps are provided for about 1/3 of all maps obtained by HIFI, with a particular emphasis on observations taken in bands 6 and 7.
- if HPDPs do not exist, the Level 2.5 spectral cubes should be used.
- irrespective of the above, in case users would like or need to re-build the spectral cubes with either improved baseline quality, or simply different re-gridding dimension, Level 2 products should be used as input for those manipulations.

Spectral Scan Observation

Due to the reasonably small number of HIFI spectral data observations ($\sim$ 500 over the whole mission), HPDPs have been created for all of them. As such we always recommend to use those as primary products for exploiting products from this observing mode. Even so, the HPDPs offer different sub-products depending on the science to be performed:

- if the users are only interested in the line information, the HPDP will provide a baseline-subtracted spectrum for each polarisation. Note that line seen in absorption in such data will not be exploitable, unless the continuum information referred to in the next bullet is used.
- in case the continuum information (alone or together with the line) is of interest, the HPDP will also contain the respective isolated continuum and the total spectrum – for this latter, the data quality is usually very similar to that of the SPG and users will most likely need to estimate a monotonous model for the continuum based on the provided inputs
- finally, in case users would like or need to re-build the deconvolved spectra with either revisited baseline correction, or potentially a different deconvolution algorithm, Level 2 products should be used as input for those manipulations.

Other products

While the above data are the prime products for the scientific analysis, a handful of other products may be of interest in order to optimise the exploitation of the data. On the one hand, the OFF spectra can inform about possible line contamination, on the other hand, line lists provide an overview of the strongest species detected in a given spectrum. Those can be queried separately as HPDPs.

6.3.2 Quality Summary

Quality Flags are raised automatically during the pipeline processing, based on event or threshold violation detection at various levels of the product generation. While part of those are mostly of interest to the instrument experts or engineers, another part brings pertinent information about possible affect to the data quality and can give hints towards the potential need for post-processing actions (Section 6.4.2).

Quality flags and associated information can be retrieved in HIPE and the HSA Science Archive in different ways and levels of information:

- The most relevant data quality information is visible in the quality summary context, i.e., in the Quality Control Summary Report. It contains both a summary (the ”comment”), as well as the list of all so-called *Public Flags* that have been passed from the pipeline-generated Quality Flags context. This summary report can be accessible through the ”quality summary” field of the QualityContext in HIPE, or alternatively in a dedicated window accessible from the HUI. Each flag comes with a description of the detected issues, an evaluation of the possible impact on science data, and an action or recommendation to follow in case the flag is raised. The summary is one of the first things, together with the postcards, that users from the HSA should inspect in the result of their data queries.
- *Private flags* contain flags being relevant only to instrument experts. They can always be inspected in the ”quality” field of the QualityContext in HIPE, but they will not be propagated to the Quality Summary context.

We have organised the HIFI quality flags into 3 categories called ”classes”, with a decreasing level of severity implied by the quality flag:

- *Class 1* flags refer to situations where the data (total or partial) is considered unusable for science
- *Class 2* flags refer to the data that is considered acceptable for science exploitation, but that in some cases could still be affected by residual instrument artefacts.
- *Class 3* flags refer to data readily usable for science with no particular further action.

A list of all quality flags possibly raised by the HIFI pipeline, together with their respective classification as of the above categories, is given in Section 10.4 of the HIFI Data Reduction Guide.

6.4 Re-processing and Post-processing of HIFI Products

We describe here specific circumstances where the HIFI data distributed as Legacy Products by the HSA could benefit from either a revision of the data processing itself from lower level products, or would need dedicated post-processing due to particular artefacts or other features present in the spectra.

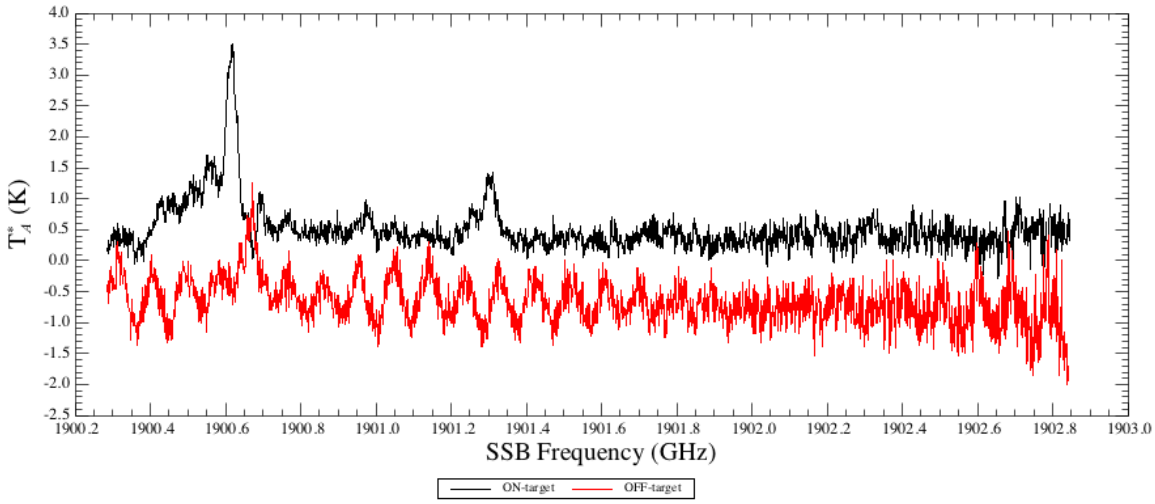

Figure 6.2: Example of line contamination in the reference position. The on-target spectrum in black shows a strong detection of the [CII] line near 1900.6 GHz, as well as a negative feature in the blue wing. This feature matches a positive emission in the isolated *OFF* spectrum shown in red, with an offset of –1 K to aid the visualisation. As expected, this latter also shows more prominent standing wave residual.

6.4.1 Re-processing HIFI Data

The HIFI products present in the HSA are the outcome of the automatic pipeline processing performed with the latest relevant version of the HIPE software (Section 6.2) and of the HIFI calibration files (Section 6.5). From a software point-of-view, this processing can be handled with a single top-level HIPE task called `hifiPipeline`. This task allows users to revisit the instrument data processing at any level, either by adjusting the input parameters to the various pipeline sub-tasks, or feeding some different calibration parameters. Because re-processing can start at any level, the use of such a task implies that lower level products (e.g. Level 0 or Level 1) need to be retrieved from the observation context. Further information about this can be found in Chapter 5 of the Data Reduction Guide.

The pipeline algorithm and its parameters have been adjusted in such a way that the best product quality is achieved for any observing circumstances. As such, a reprocessing of the HIFI data should be considered very cautiously. There are, however, specific circumstances where the default pipeline could not reach the best possible science output – these situation are usually very linked to the source properties and the (in-)adequacy of the observing mode to optimally extract the relevant scientific information. We list thereafter the most representative cases. It should be noted that the recipes provided thereafter assume the usage of HIPE and its tasks. Even if users wish to use other software packages to deal with the highlighted data issues, we recommend to read the following to be familiar with.

Contamination from the OFF position

There exist rare cases where the data acquired on the reference position was not strictly free of emission, typically because the reference position was chosen to be too close to the main source, or when

the environment in the source surrounding was very crowded. In calibrated data this will manifest as spectral artefacts appearing as negative features (although they can also be positive in certain cases of Double Beam Switching observations – Section 3.3.1). These are referred to as *OFF contamination*, and they can lead to misinterpretation of line strength and profile, but also of the real continuum level in case the reference position also features non-negligible continuum levels. Figure 6.2 illustrate such a situation in the case of a [CII] line detected towards a Galactic Plane position.

The cleanliness of the OFF position can be assessed by inspecting the isolated OFF spectra also generated by the standard pipeline – the various cookbooks provided in Section 2 of the Data Reduction Guide indicate where these can be found, and how they are computed for each observing mode. Those OFF spectra are also provided as a dedicated Highly-Processed Data Product (Section 6.2.3).

There are two possible ways to correct for those contamination. In most cases, however, it should be noted that these corrections will lead to enhanced baseline artefacts since the OFF subtraction was here to cancel them out in the first place.

- for observations using a double referencing scheme (Section 3.3), namely DBS, Load Chop with Reference, and Frequency Switching with Reference, it might actually be more advantageous to skip the subtraction of the reference position, therefore ignoring the term $J_{\mathrm{OFF,ssb}}$ in Eq. 4.1. This requires an *ad hoc* reprocessing of the HIFI data whereby the task `doOffSubtract` should simply be de-activated in the configurable pipeline - see Section 5.4.2 of the Data Reduction Guide. Note that this option is particularly interesting in bands 6 and 7 where the main drift component comes from the Electrical Standing Wave (see Section 5.3), which is optimally corrected following a technique that does not necessarily require the subtraction of the reference position.

- alternatively, and irrespective of the mode, users can also directly use the isolated OFF spectrum and add it to the top-level ON-target spectrum. Since the isolated OFF data usually suffer from baseline artefacts, it is recommendable to first try and clean those (following the same recipes as given e.g. in Section 6.4.2) before combining them to the ON data. For Spectral Scans in particular, it may be interesting to correct for the whole deconvolved OFF spectrum, which is also provided as Legacy Product and HPDP (Section 6.2).

Whenever an OFF contamination could be identified by one of the instrument expert, a dedicated Quality Report entry is provided in the Quality Summary of the product. For Spectral Maps provided as HPDPs, the OFF contamination is corrected in some cases (see Morris 2016).

Residual Electrical Standing Waves

The vast majority of the ESW affecting the HIFI data from bands 6 and 7 are well corrected through the standard pipeline. In rare cases, however, this correction is not yet optimal, usually due to more complicated line profiles that were not properly masked by the automatic line identifier. In this case, it is possible to inform the ESW correction task `doHebCorrection` about the adequate line masking window and re-run the ESW correction. This implies working from the Level 1 products, and can be significantly CPU- and memory-consumming for large datasets (especially maps). Recipes on how to perform this reprocessing can be found in Section 12.4 of the Data Reduction Guide.

Re-running the sideband deconvolution in Spectral Scans

Although Legacy Products associated to Spectral Scans have gone through an extensive spur and baseline artefact cleaning (also delivered as HPDP, Section 6.2.3), users might want to re-run the sideband deconvolution on an improved version of the Level 2 individual spectra. Once the re-work on Level 2 data is complete, information on how to re-run the deconvolution task can be found in Chapter 14 of the Data Reduction Guide.

Strong continuum sources and enhanced optical standing waves

HIFI data users should pay particular attention to observations of relatively large (typically higher than 5 K in T_A^*, or $\simeq$ 2500 Jy) continuum levels – this is particularly true for planets such as Mars or Jupiter. Indeed, because the modulation of residual optical standing wave is multiplicative, their magnitude will be exacerbated in the case of large source continuum. Figure 6.3 shows the extreme case of an observation of Mars in band 1b, near the HDO line at 600 GHz. In this example, a continuum of about 10 K is observed, giving rise to a standing wave oscillation of magnitude $\simeq$ 0.2 K (it can actually be seen that this oscillation has an amplitude modulation due to the beat of two very close optical resonances – those from the hot and cold load cavities respectively – see also Section 5.3 and Rodríguez-Coira & Teyssier 2015). Since standing waves are present both in the Lower and Upper Sidebands, the final spectrum is made of the superposition of the two. Because they may not necessarily add up with the same phase, this superposition can lead to constructive or destructive oscillations, making the overall modulation worse or better. The case illustrated here for example is unusually bad, and shows how the H polarisation data suffers more than the V polarisation data in that respect.

Because the resulting oscillation is intrinsically due to the fact that a modulation present in the hot and cold load counts gets multiplied by the counts from the on-target phase, one option consists in mitigating the modulation in the load counts prior to their multiplication (Risacher 2011). An option exists in the configurable pipeline to perform this mitigation, albeit at the expense of other possible knock-on effects. It is described in more detail in Section 12.2 of the Data Reduction Guide.

6.4.2 Post-Processing of HIFI Data

As explained in Section 6.3.1, there are two main kinds of post-processing steps that may be required on the HIFI archival data. The first one stems from the particularity of the intensity calibration scheme and relates to transformations needed on the intensity scales, both for the line and continuum information. The second one is more directly linked to the items highlighted in Section 5.3 and deal with the necessary data cleaning that may still apply to some of the Level 2 or Level 2.5 products. Part of these latter caveats will be resolved in the framework of the HPDPs (Section 6.2), however, this will not apply to all products potentially still affected by artefacts. We briefly describe here the typical workflows associated with either of the circumstances. These details are also provided on an observing-mode-basis in all the cookbooks present in the HIFI Data Reduction Guide.

Prologue: Polarisation averaging

Irrespective of the items introduced above, one preliminary step that users should consider doing whenever possible is the combination of the data collected in both polarisations in order to beat down

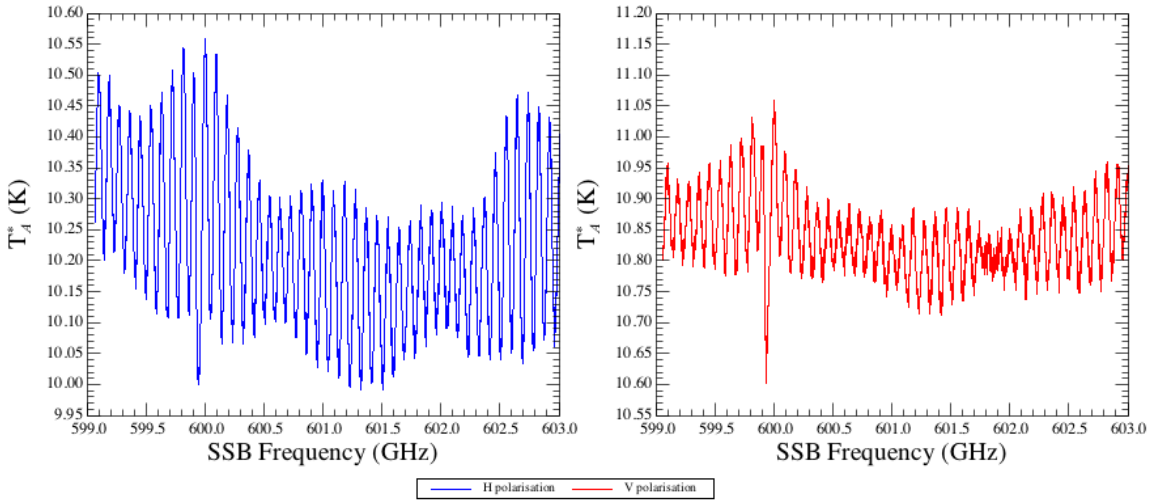

Figure 6.3: Illustration of enhanced residual standing wave amplitude in the case of a strong continuum source (here Mars taken from obsid 1342194498).

the noise by up to a factor $\sqrt{2}$. There are precautions to take when doing this, the background of it being given in Section 5.8.4. Also, for maps, it is recommended to do this merging at the time of the cube generation – see also Section 15.3.2 of the Data Reduction Guide.

Residual artefacts: Standing Waves and Baseline drifts

The wish of any heterodyne user is to be able to work on flat baseline data. Whenever this is not already the case in the final products, baseline distortion will need to be further removed in order to estimate accurate line intensities. In HIPE, this can be done with the respective tasks *fitBaseline* and *fitHifiFringe*, however, this can obviously be done in any of the users' favourite software.

In Spectral Scan and Spectral Map data, this correction needs to be performed in the Level 2 data, then the respective deconvolution and re-gridding steps need to re-run with the cleaned data ingredients. Many of those observations may actually have already benefited from such treatment as HPDPs so you should check whether these are already available for your observation of interest.

A summary is given next on the presence and correction methods for residual standing waves in Level 2 data that have all been reference-corrected in the standard fashion.

Bands 1-5: DBS Modes

Bands 1-5 data taken in DBS mode generally show very little standing wave residuals at Level 2. Exceptions are sometimes observed, however, and in these cases the user can remove the waves with the sine wave fitting task `fitHifiFringe` in HIPE. Two cases are shown below.

- In Band 4b the diplexer causes a 650 MHz ripple. Figure 6.4 shows a sine wave fit (red) relative to a baseline (green), automatically determined with *fitHifiFringe*.

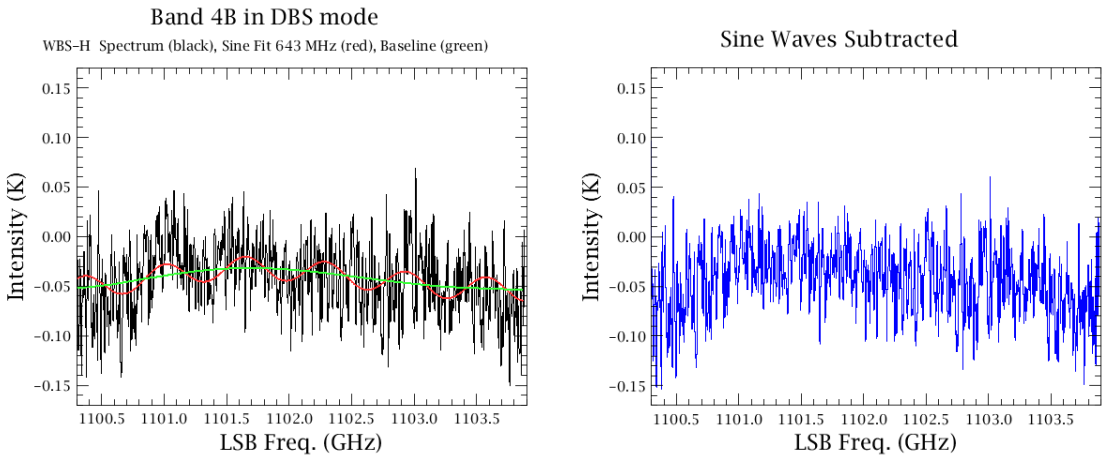

Figure 6.4: Standing wave residual at Level 2 in a DBS observation on the left, with the fitted sine wave (red) relative to a smoothed baseline (green), and corrected spectrum on the right.

- Strong continuum sources show the effect of ripples in the passband calibration. The periods are typically in the 90-100 MHz range, and the amplitudes are at most 2% relative to the continuum. See Figure 6.5 for a case where 92 and 98 MHz sine waves are present.

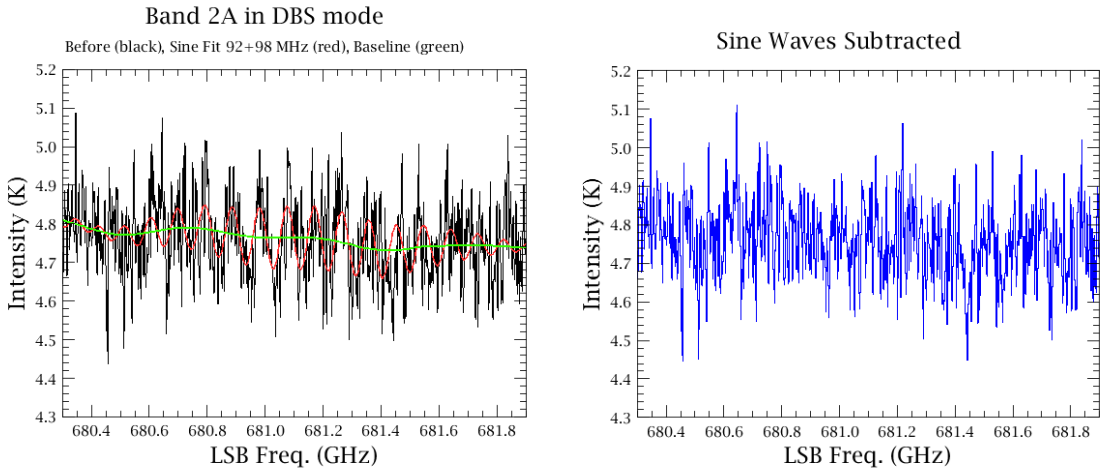

Figure 6.5: Strong continuum source with two standing wave components (red) relative to the baseline (green) on the left, and corrected spectrum on the right.

Bands 1-5: Position Switch, Frequency Switch, Load Chop Modes (Point and Spectral Scan AOTs)

Like calibrated observations taken with the Point or Spectral Scan AOTs using the DBS modes, the Position Switch, Frequency Switch (with sky reference measurement), or Load Chop (with sky reference measurement) modes in beamsplitter Bands 1, 2, and 5 also show the cleanest calibrated spectra.

If any residual waves are present in the Level 2 products, they have the shape of pure sine waves and can be subtracted using the *fitHifiFringe* task in HIPE.

Position-switched observations in diplexer Bands 3 and 4 often show residual waves with larger amplitudes compared to the DBS modes. Waves generated in the diplexer rooftop are not pure sine waves. Amplitudes increase strongly toward the IF band edges. It was thus always advised to place the line of interest near the middle of the IF band when possible, in WBS sub-bands 2 or 3. Although *fitHifiFringe* only fits sine waves, fits to the diplexer waves can be approximated using multiple sine waves, typically around 600 MHz. The approximation is not as good at the IF band edges. Examples are shown for Bands 3 and 4 in Figure 6.6 respectively. Note that the baseline is strongly curved in both observations, and the user can treat those separately in HIPE, using polynomials with `fitBaseline`.

Bands 1-5: Mapping Modes

Refer to the discussion on baseline drift and standing wave effects on the mapping modes in Section 5.5.3.

Bands 6-7: DBS Modes

The HEB mixers generate waves with characteristics that depend on the mixer current level, but since there is a dependence on LO power which may not be stabilised (between ON and OFF spectra), both electrical and optical standing wave residuals can be mixed together. Data taken with all observing modes analysed so far show residuals at Level 2, where the non-DBS modes tend to show larger and more complex residuals.

Bands 6-7 are less stable than Bands 1-5, and Level 2 data occasionally show residual waves even in the chopped modes. The ESWs have periods of $\sim$300 MHz, but they are not sine waves. HIFI users were given the general guidelines that: Band 6b is most stable and least affected by residual ESWs; and FastDBS mode spectra tend to show weaker standing waves than 'slow chop' DBS mode spectra.

Bands 6-7: Position Switch, Frequency Switch, Load Chop Modes

Before the introduction of the `hebCorrection` task in the HIFI pipeline (ca. SPG v12.1), non-DBS observations could show rather strong residuals in Level 2 spectra, with no specific tendency for the strengths or number of affected observations to be any better or worse by mode. However, since the problems occur primarily at frequencies where LO power has long stabilisation times, we might expect Frequency Switch observations to have the poorest quality baselines with the artefacts of drifting instrumental response between measurements at the two frequency phases. That there are ugly Frequency Switch data is true, but there are no significantly fewer affected observations using standard Position Switch or Load Chop.

The best method to deal with ripples in the HEB bands has proven to be with the `hebCorrection` task, as described in Section 6.4.1.

Residual artefacts: Deconvolution of strong lines

The deconvolution algorithm can create spectral ghost features (usually showing up as negative features) in the neighbourhood of relatively strong lines (typically several tens of Kelvins). This is illustrated in Figure 6.7. In order to recover both weak and strong spectral line emission in the Level 2.5 SSB spectrum, it is necessary to mask those strong lines in the Level 2, and then run the deconvolution in two steps, with the strong lines unmasked, and with the strong lines masked.

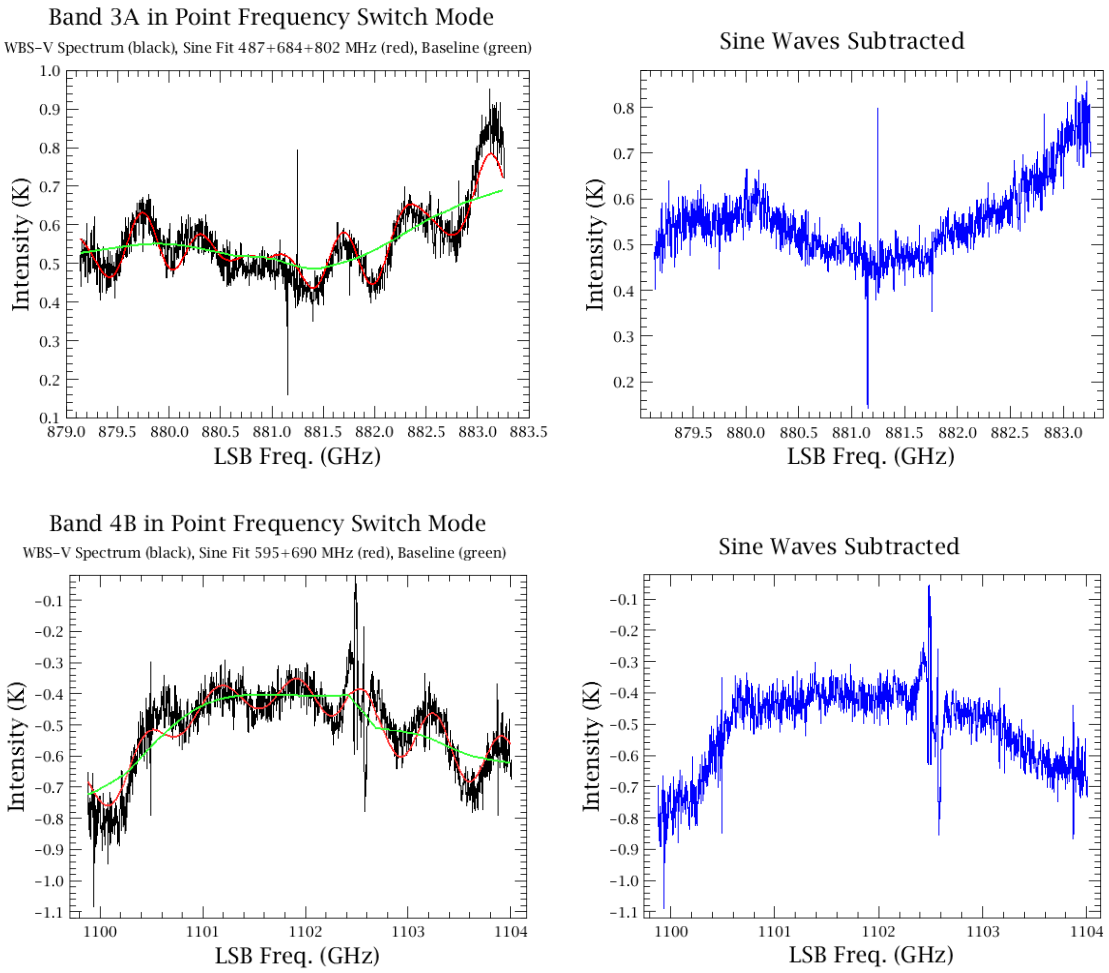

Figure 6.6: *Upper panels*: Band 3a observation using Point Mode Frequency Switching at Level 2 after sky subtraction (left) and after interactive fringe fitting and removal (right). Note the poor edges in the spectrum on the right. *Lower panels*: Band 4b observation using Point Mode Frequency Switching, showing a strong residual standing wave (black). Its period (∼600 MHz) and increasing amplitude toward the band edges are those expected for diplexer bands. It was fitted with a combination of sine waves using *fitHifiFringe* (green). The corrected spectrum (blue) still shows a strongly curved baseline that can be subtracted with a polynomial fit in HIPE.

In the above result where the strong lines were masked, the bright line will be totally absent from the SSB spectrum, but so will be the negative ghosts. It is then a matter of merging the two above deconvolution outputs in order to basically inject the isolated signal of the bright lines recovered from deconvolution applied with the bright lines unmasked into the spectrum from the deconvolution with the strong lines flagged. The Spectral Scan cookbook gives further recipe and reference to a useful script allowing to perform this correction.

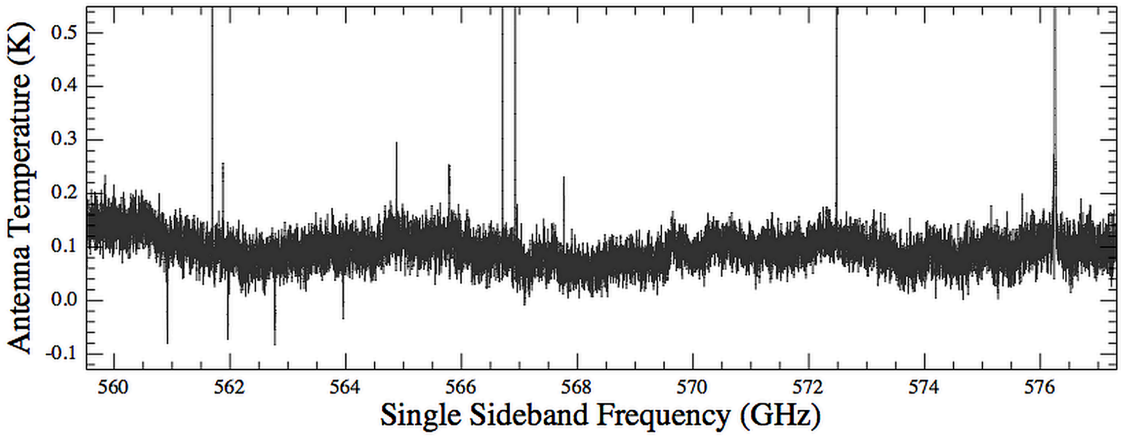

Figure 6.7: Spectral Ghost artefacts from the deconvolution of a strong line (standard Level 2.5 products from the WBS-H in Obsid 1342215923). The strong ^{12}CO line lies at 576.5 GHz (70 K), and injects negative signal in the deconvolution in the range 561–564 GHz.

Residual artefacts: Deconvolution of Spectral Scans in Frequency Switching mode

Taking Spectral Scan in combination with the FSW referencing scheme was essentially limited to the Performance Verification phase, and in effect, almost no science program did finally use it over the mission. In total, there are 19 Spectral Scans taken in this fashion in the Herschel Science Archive (only four are not from the Calibration programme). In general, the data quality of these data meets the general level achievable with other referencing schemes (with the exception of bands 6 and 7 where it does not do a good job at all), however there can be issues when it comes to deconvolving those data.

We recommend using these data and the outcome of their deconvolution with very special care, and favour any equivalent dataset taken with another referencing scheme such as DBS or Load Chop.

Epilogue: Intensity scale transformation

We explain in Section 4.1.1 why the T_A^* antenna temperature scale applied to the Level 2 and Level 2.5 products is not readily usable to perform scientific analysis of the HIFI data. The possible scale and/or unit conversion steps are explained in detail in Section 5.8.6. In case both continuum and line information need to be exploited, additional data manipulation is needed, as explained in Section 5.8.5.

6.5 Access to Calibration Data: HIFI Calibration Tree

Each HIFI observation is accompanied by a complete set of calibration file for that observation. The calibration parameters can be both general (applicable to all observations) or specific to that observation. This section provides more details about the various items within the calibration files (also called

calibration tree). The calibration files are divided into three main components:*Uplink*, *Downlink*, and *pipeline-out*.

Information about how the observation was planned are contained in the *Uplink* node of the calibration tree and listed in Table 6.4. Information relevant to the downlink processing are contained in the *Downlink* node of the calibration tree and listed in Table 6.5. Downlink calibration parameters tend to be global calibration files, however, there are a few which are observation specific. Finally, information computed in the course of pipeline processing is under the *pipeline-out* node and listed in Table 6.6.

Each component can be further divided into sub nodes. Sub-nodes can refer to specific pipeline processing stages (e.g., Level0) or specific processing steps. The list provided here does not list all parameters available within the HIFI calibration tree, just those which may aid in future re-processing or for deeper understanding of the observation.

Product	Description
HifiAORData	
the Obsid value	If present, indicates the values and parameters used in HSPOT to plan the observation
HifiUplinkData	
HifiUplinkModes	List of All HIFI modes. Note this is observation independent
HifiUplinkParameters	List of parameters which belong to various uplink modes (also not observation dependent)

Table 6.4: Calibration Uplink Parameters

Product	Description
Eng	
DiplexerCoefficients	Table containing the diplexer coefficients for the diplexer bands (3, 4, 6 and 7)
Generic	
BeamProfiles	Tables containing the HIFI beam profiles. There is a node containing the tables for the 2D beam model and a node containing the tables for the azimuthally-averaged (1D) beam models. Beam models are provided for two spot frequencies, per mixer band for each polarisation H and V. Spot frequencies are chosen to be close to the central frequency of the LO subbands a and b, respectively (Section 5.9.2, Mueller et al. 2014)
Spurs	**Obsolete**: historically, tables listing known LO impurities. Now replaced by spurFlags
apertureEfficiency	Tables listing the aperture efficiencies for each band for the H and V polarisations, respectively
beamEfficiency	Tables listing beam efficiencies for each band for the H and V spectrometers, respectively.
beamWidth	Tables listing beam widths for each band for the H and V polarisations, respectively.
chopperPositions	Table listing the voltages required to point the chopper at each load. There is one table for the prime (not used after OD 81) and redundant sides.
chopperThrows	Table listing the chopper throws (in degrees) for each band

couplingEfficiency	Tables listing the efficiency in coupling to the hot and cold loads for each band for the H and V polarisations, respectively.
forwardEfficiency	Tables listing the forward efficiency for each band for the H and V polarisations, respectively.
hebCorrection	Spline models and fit results of standing waves found in the HEB mixers (bands 6 and 7). These are also known as ESW (Section 5.3.1).
mixerCurrentTolerances	Table listing the mixer current tolerances for each band
sideBandGainIF	Tables listing the sideband gain coefficients in the IF for each band for the H and V polarisations, respectively.
sideBandGainLO	Tables listing the sideband gain coefficients in the LO for each band for the H and V polarisations, respectively.
smoothOffWidth	Tables listing the smoothing widths in MHz, per band, to be used in the pipeline for Load Chop and Frequency Switch observations for the H and V polarisations, respectively.
spurFlags	List of spurs to be applied as channel and rowflags to a given Obsid. For point and map observations, the vast majority is based on a spur predictor model and labelled as `SPUR_WARNING` (de Jong & Teyssier 2016)
uncertaintyModel	Values for uncertainty parameters described in Section 5.8.2
High Resolution Spectrometers	
CalHrsPowCorr	Gain non-linearity correction table used in the HRS pipeline
CalHrsQDCFast	Table needed to perform an approximate correction of the quantisation distortion of the correlation functions of HRS. This table is not used as a standard part of the HRS pipeline
CalHrsQDCFull	Table needed to correct the quantisation distortion of the correlation functions of HRS. This calibration is used in standard processing of HRS data
Level 0	
APE	Measured absolute pointing error for different time periods of Herschel operations.
BBids	Table listing BBids used in the pipeline to identify observing mode and phases of the observation
CleanDF	Tables of observations (obsids) which were found to be corrupted (Section B)
HK	Table listing house-keeping parameters, their units and a description.
UpConvertLo	The upconversion factors (in the metadata)
hkThreshold	Tables listing the housekeeping threshold values for the H and V chains, respectively, for: the mixer currents of the SIS bands (1-5), the magnetic resistance for all bands; the diplexer current for the diplexer bands (3, 4, 6 and 7); and the Local Oscillator Unit (LOU) currents for bands 1, 3, and 7 (to avoid impurities)
pipelineParameter	An Obsid-specifc table indicating where parameters of a given pipeline task should differ from their default
Wideband Spectrometers	
badPixels	Table listing the pixels known to be bad for the WBS-H and WBS-V polarisations, respectively.
combFitParameters	
linearityCoefficient	Table listing the WBS linearity coefficients for each band for the WBS-H and WBS-V polarisations, respectively.

zeroThresholds	Table of minimum and maximum acceptable Zero values for the WBS-H and WBS-V polarisations, respectively.

Table 6.5: Calibration Downlink Parameters

Product	Description
Generic	
ESWCorrection	The list of fit parameters per spectrometer of the found ESW for each spectrometer (Section 5.3.1).
Baseline	Only for observations containing a sky reference (OFF) position. Spectra of the smoothed and averaged OFF positions for each of the spectrometers.
ReferenceSpectra	Calibrated sky reference (OFF) spectra. For DBS observations, these are difference spectra of the averaged CHOP positions for each spectrometer.
FrequencyGroups	Tables listing the dataset index mapping for the frequency groups identified by the pipeline for each of the spectrometers.
Tsys	Spectra of the Tsys and BandPass calculated by the pipeline for each of the spectrometers.
Uncertainty	Individual and combined uncertainty parameters per spectrometer and polarisation. For HRS, subbands are listed separately. For WBS, uncertainties at low, mid and high IF frequencies are indicated.
WBS	
BadPixelProposed	Table listing pixels identified by the WBS pipeline to be saturated for the WBS-H and WBS-V spectrometers, respectively.
WbsFreq	Frequency coefficients calculated by the MkWbsFreq step of the WBS pipeline for the WBS-H and WBS-V spectrometers, respectively.
Zero	Spectra of the zero level observations for the WBS-H and WBS-V spectrometers, respectively.

Table 6.6: Calibration Pipeline-out Parameters

6.6 Working with HIFI data in CLASS/Gildas

CLASS, which is part of the GILDAS package[1], is a very commonly used software tool for the analysis of single dish radio-astronomy data. At the beginning of the mission, HIFI data needed a special conversion to the Class data format (through the HIPE `HiClass` task), however, this is no longer needed and CLASS now reads HIFI FITS files from the HSA, *provided you use the October 2015 version or more recent*. The products covered by this importer are both the HIFI Level 2 (HIFI Timeline Products) and Level 2.5 products. In this process, a portion of the HIFI header information is lost, however, all fundamental HIFI House-Keeping parameters have been migrated to either the standard Class data header, or alternatively, to a newly introduced variable section called `R%HEAD%HER%`, that contains information specific to HIFI. Noticeable additions are for example those covering the cou-

[1] https://www.iram.fr/IRAMFR/GILDAS/

pling efficiencies and beam properties. Details about the importer and the new variable section can be found in Bardeau et al. (2015).

In the following, we provide some basic Class command instructions needed to import that HIFI data. It is assumed here that users are already familiar with the CLASS environment and operations. Further information can be found in the CLASS documentation[2]. We want to warn that all commands and URLs used here work at the time of writing but the HSC cannot guarantee that CLASS commands will remain with the current syntax.

The script below shows how Level 2.5 products downloaded from the HSA can be simply imported into CLASS. We are dealing with the standalone browse products delivered for obsid 1342191638, here restricted to the WBS ones. The respective sideband spectra are stored in two newly created files lsb1113.hifi and usb1113.hifi.

```
LAS> file out lsb1113.hifi s /o
LAS> fits read hhifiwbshlsb1342191638_25ssv20_1452784743777.fits
LAS> fits read hhifiwbsvlsb1342191638_25ssv20_1452784743568.fits
LAS> file out usb1113.hifi s /o
LAS> fits read hhifiwbshusb1342191638_25ssv20_1452784743568.fits
LAS> fits read hhifiwbsvusb1342191638_25ssv20_1452784743568.fits
```

At this stage, the user will find the data in the same shape as they were in the HSA. As an example, the frequency scale will be on the LSR scale, and so for example the following would be needed to get an accurate velocity information for the line of interest:

```
LAS> get f !  read your spectrum into memory
LAS> modify frequency 1113342.964 !  water ground state, as an example
```

Similarly, if you need to transform the T_A^* scale into a main temperature scale, the following can be done:

```
LAS> modify beam_eff R%HEAD%HER%ETAMB ! should contain appropriate
efficiency from Table 5.8
```

Import of HIFI channel flags

One of the powerful information in the HIFI products are the various flag categories associated to given channels. Those flags are in particular essential in order to perform proper sideband deconvolution in Spectral Scan observations, so that running the equivalent deconvolution in CLASS would fail if those flags would not be ported to the CLASS data. To this end, a new feature has been introduced in the October 2015 CLASS version, that allows to maintain this flag information into the data. These new *associated arrays* (Bardeau & Pety 2015) currently allow to map the HIFI flags about artefacts (typically `BAD_DATA` and `IGNORE_DATA`) and about lines (typically `LINE` and `BRIGHT_LINE`) into dedicated variables, which can then be used to adequately mask the data:

- the BLANKED Associated Array is an integer flag array indicating if the RY values (i.e. the spectrum) should be blanked (1) or not (0). CLASS generally uses the value -1000 for bad

[2] https://www.iram.fr/IRAMFR/GILDAS/doc/html/class-html/class.html

channels. The HIFI spectra keep their original value and the flag array is included. Generally, the flagging is appropriate and the corresponding masking can be performed:

```
LAS> LET RY R%HEAD%SPE%BAD /WHERE R%ASSOC%BLANKED%DATA.EQ.1
```

- The channels that have been identified as containing lines are stored in the array R%ASSOC%LINE%DATA and they can be used to define baseline masking windows as follows:

```
LAS> set window /associated !  puts the line windows into the arrays
w1 and w2
LAS> base 1 /pl !  subtracts and plots a linear baseline using the
HIPE defined windows
```

Note that these arrays are currently lost when data are smoothed.

What important metadata will be missing ?

The LO frequency is only contained in the line name and is no longer numerical. However, the signal and image frequencies are present.
Also, since the velocity frame for Solar System Objects is linked to the object and not LSR, the type is set to *unknown* in CLASS as this frame is not recognised there.

Chapter 7

Bibliography

Avruch, I. M. 2011, *Cycle38 Frequency Calibration Tests*, Tech. Rep. HIFI-ICC-TN-2015-002, SRON Groningen

Avruch, I. M. & Melchior, M. 2011, *Frequency and Velocity Transforms for HIFI*, Tech. Rep. ICC/2011-003, SRON Groningen

Bardeau, S. & Pety, J. 2015, The new CLASS Associated Arrays, Memo 2015-4, IRAM

Bardeau, S., Teyssier, D., Rabois, D., & Pety, J. 2015, Importing Herschel-FITS into CLASS, Memo 2015-3, IRAM

Belgacem, M., Ravera, L., Caux, E., Caïs, P., & Cros, A. 2004, New A, 9, 43

Benamati, L. & Caux, E. 2016, in preparation

Bjerkeli, P., Liseau, R., Larsson, B., et al. 2012, A&A, 546, A29

Cherednichenko, S., Drakinskiy, V., & Berg, T. 2005, in Proc. Europ. Microwave Association, European Microwave Association

Cherednichenko, S., Kroug, M., Merkel, H., et al. 2002, Physica C Superconductivity, 372, 427

Comito, C. & Schilke, P. 2002, A&A, 395, 357

Crockett, N. R., Bergin, E. A., Neill, J. L., et al. 2014, ApJ, 787, 112

de Graauw, T., Helmich, F. P., Phillips, T. G., et al. 2010, A&A, 518, L6

de Graauw, T., Whyborn, N., Helmich, F., et al. 2008, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7010, Space Telescopes and Instrumentation 2008: Optical, Infrared, and Millimeter, 701004

de Jong, A. & Teyssier, D. 2016, Spur Predictor for the HIFI point and mapping observations, Tech. rep., SRON-G

De Luca, M., Gupta, H., Neufeld, D., et al. 2012, ApJ, 751, L37

Delorme, Y., Salez, M., Lecomte, B., et al. 2005, in Sixteenth International Symposium on Space Terahertz Technology, 444–448

Dieleman, P. 2009, HIFI SVM panel thermal stability report, Tech. rep., SRON-G

Emerson, D. T. & Graeve, R. 1988, A&A, 190, 353

Feuchtgruber, H. 2012, Reconstruction of the Herschel Pointing Jitter, Tech. rep., MPE

Gerin, M., de Luca, M., Goicoechea, J. R., et al. 2010, A&A, 521, L16

Goicoechea, J. R., Teyssier, D., Etxaluze, M., et al. 2015, ApJ, 812, 75

Goldsmith, P. F., Yıldız, U. A., Langer, W. D., & Pineda, J. L. 2015, ApJ, 814, 133

Griffin, M. J., Abergel, A., Abreu, A., et al. 2010, A&A, 518, L3

Harris, A. I. 1988, International Journal of Infrared and Millimeter Waves, 9, 231

Hartogh, P., Lellouch, E., Crovisier, J., et al. 2009, Planet. Space Sci., 57, 1596

Herpin, F. & Teyssier, D. 2003, Frequency calibration framework document, Tech. rep., SRON-G

HIFI-ICC. 2013, Release notes to the HIFI observations taken in the framework of the Performance Validation campaigns, Tech. rep., SRON

Higgins, D. R. 2011, PhD thesis, National University of Ireland Maynooth

Higgins, R. & Kooi, J. 2009, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7215, Terahertz Technology and Applications II, ed. K. Linden, L. Sadwick, & C. O'Sullivan

Higgins, R., Teyssier, D., Borys, C., et al. 2014, Experimental Astronomy, 37, 433

Jackson, B., de Lange, G., Zijlstra, T., et al. 2006, Microwave Theory and Techniques, IEEE Transactions on, 54, 547

Jackson, B. D., de Lange, G., Zijlstra, T., et al. 2005, Journal of Applied Physics, 97

Jackson, B. D., Wildeman, K. J., & Whyborn, N. D. 2002, in Thirteenth International Symposium on Space Terahertz Technology, 339

Jellema, W. 2008, HIFI Spacecraft Instrument Alignment Matrix, Tech. Rep. ICC/2008-130, SRON Groningen

Jellema, W. 2015, PhD thesis, SRON-G

Jellema, W., Jacobs, H., van Leeuwen, B.-J., et al. 2010, in Twenty-First International Symposium on Space Terahertz Technology, 33–39

Jellema, W., Jochemsen, M., Peacocke, T., et al. 2008, in Ninteenth International Symposium on Space Terahertz Technology, ed. W. Wild, 448

Kester, D., Avruch, I., & Teyssier, D. 2014, Bayesian Inference and Maximum Entropy Methods in Science and Engineering, 1636, 62

Kester, D., Higgins, R., & Teyssier, D. 2016, in press

Kooi, J. 2009, HIFI IF Stability on OD16 and the thermal effect of the SVM panels, Tech. rep., SRON-G

Kooi, J. & V., O. 2009a, HIFI Differential Instrument Stability as measured during the CoP phase, Tech. rep., SRON-G

Kooi, J. & V., O. 2009b, HIFI System Stability Performance – COP, Tech. rep., SRON-G

Kramer, C. 2008, Spatial response framework document, Tech. Rep. SRON-G/HIFI/TR/2008-010, SRON Groningen

Kutner, M. L. & Ulich, B. L. 1981, ApJ, 250, 341

Larsson, B. & Teyssier, D. 2008, HIFI FM Radiometry Measurements, Tech. Rep. SRON-G/HIFI/TR/2008-001, SRON Groningen

Morris, P. 2011, HIFI AOT Observing Mode Release and Performance Notes, Tech. rep., SRON

Morris, P. 2015, Pointing Reconstruction in HIFI Observations, Tech. rep., NHSC

Morris, P. 2016, Herschel/HIFI Spectral Map Highly Processed Data Product Release Notes, Tech. Rep. HERSCHEL-HSC-DOC-2122, ESAC

Morris, P. & Ossenkopf, V. 2017, in preparation

Moustakas, L. A., Cyr-Racine, F., & Keeton, C. R. 2014, in American Astronomical Society Meeting Abstracts, Vol. 223, American Astronomical Society Meeting Abstracts #223, 408.04

Mueller, M., Jellema, W., Olberg, M., Moreno, R., & Teyssier, D. 2014, *The HIFI Beam: Release 1 – Release Note for Astronomers*, Tech. Rep. HIFI-ICC-RP-2014-001, SRON Groningen

Müller, T., Balog, Z., Nielbock, M., et al. 2014, Experimental Astronomy, 37, 253

Neufeld, D. A., Sonnentrucker, P., Phillips, T. G., et al. 2010, A&A, 518, L108

Nielbock, M. 2010, PACS pointing calibration sources, Tech. rep., MPIA

Olberg, M. 2015, Routine phase DBS raster maps, Tech. rep., OSO

Olberg, M. 2017, in preparation

Ossenkopf, V. 2003, The HIFI intensity calibration framework, Tech. Rep. 442, University of Cologne

Ossenkopf, V. 2008, A&A, 479, 915

Ossenkopf, V. 2015, The propagation of systematic uncertainties in the intensity calibration for HIFI, Tech. rep., University of Cologne

Ossenkopf, V. & Morris, P. 2005, HIFI Observing Mode Description, Tech. rep., SRON

Pearson, J. C., Guesten, R., Klein, T., & Whyborn, N. D. 2000, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 4013, UV, Optical, and IR Space Telescopes and Instruments, ed. J. B. Breckinridge & P. Jakobsen, 264–274

Pilbratt, G. L., Riedinger, J. R., Passvogel, T., et al. 2010, A&A, 518, L1

Poglitsch, A., Waelkens, C., Geis, N., et al. 2010, A&A, 518, L2

Rhoads, J. E., Malhotra, S., Allam, S., et al. 2014, ApJ, 787, 8

Risacher, C. 2011, Calibration impact of HIFI optical Standing Waves, Tech. rep., SRON-G

Rodríguez-Coira, G. & Teyssier, D. 2015, Investigation and correction of Herschel/HIFI spectral baseline artefacts, Tech. rep., ESA-ESAC

Roelfsema, P. R., Helmich, F. P., Teyssier, D., et al. 2012, A&A, 537, A17

Rudy, D. J., Muhleman, D. O., Berge, G. L., Jakosky, B. M., & Christensen, P. R. 1987, Icarus, 71, 159

Salez, M., Delorme, Y., Peron, I., et al. 2003, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 4855, Millimeter and Submillimeter Detectors for Astronomy, ed. T. G. Phillips & J. Zmuidzinas, 402–414

Sánchez-Portal, M., Marston, A., Altieri, B., et al. 2014, Experimental Astronomy, 37, 453

Schieder, R. & Kramer, C. 2001, A&A, 373, 746

Schieder, R. T., Siebertz, O., Gal, C., et al. 2003, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 4855, Millimeter and Submillimeter Detectors for Astronomy, ed. T. G. Phillips & J. Zmuidzinas, 290–300

Shipman, R. 2017, in preparation

Teyssier, D. 2016a, HIFI Beams as Ancillary Data Products: Release notes, Tech. Rep. HERSCHEL-HSC-DOC-2112, ESAC

Teyssier, D. 2016b, in preparation

Teyssier, D. 2016c, HIFI Non-Averaged Level 2 products for Point Mode Observations: Release notes, Tech. Rep. HERSCHEL-HSC-DOC-2103, ESAC

Teyssier, D. 2016d, HIFI Reference Position Spectra Data Products: Release notes, Tech. Rep. HERSCHEL-HSC-DOC-2111, ESAC

Teyssier, D. 2016e, HIFI System Noise Temperature Data Products: Release notes, Tech. Rep. HERSCHEL-HSC-DOC-2113, ESAC

Teyssier, D. 2016f, HIFI Uncertainty Table Data Products: Release notes, Tech. Rep. HERSCHEL-HSC-DOC-2110, ESAC

Teyssier, D. 2017a, in preparation

Teyssier, D. 2017b, in preparation

Teyssier, D. 2017c, in preparation

Teyssier, D. & Edwards, K. 2016, HIFI Trend Analysis Data Products: Release notes, Tech. Rep. HERSCHEL-HSC-DOC-2104, ESAC

Teyssier, D. & Higgins, R. 2013, Side-Band Ratio Correction for HIFI Data, Tech. rep., ESAC

Teyssier, D. & Osenkopf, V. 2017, in preparation

Van Dyk, S., Morris, P., & Jones, S. 2016, Empirical Examination of Possible H/V Imbalance, Tech. rep., NHSC

Ward, J., Maiwald, F., Chattopadhyay, G., et al. 2003, in Fourteenth International Symposium on Space Terahertz Technology, ed. C. Walker & J. Payne, 94

Wilson, T. L., Rohlfs, K., & Hüttemeister, S. 2009, Tools of Radio Astronomy (Springer Berlin Heidelberg)

Appendices

Appendix A

The HIFI SIAM

The Spacecraft-Instrument Alignment Matrix (SIAM) is a collection of rotation matrices, each of which describes the necessary transformation needed to point the telescope along one of the possible instrument line of sights. These rotation matrices are defined by the offsets between an instrument aperture in the focal plane and the coordinate system used by the Attitude Control and Measurement System (ACMS).

The total number of apertures required by HIFI is defined as follows (Jellema 2008):

- 7 independent mixer bands
- in each band there are 2 LO subbands (frequency dependence)
- in each band there are 2 mixers operating on orthogonal polarisations (polarisation dependence)
- in each band there are 2 relevant chopper positions (chopper dependence)

For each mixer band we consequently include frequency (2 sub-entries), polarization (2 sub-entries) and chopper (2 sub-entries) dependence in the SIAM. The total number of apertures included in the SIAM for HIFI is therefore $7 \times 2 \times 2 \times 2 = 7 \times 8 = 56$ entries.

Frequency dependence

It is verified by analysis that the HIFI beams are not frequency dependent as far as pointing is concerned. However, in each mixer band there are two LO subbands designated by “A” and “B” that subdivide the mixer bands into two LO frequency subbands. Since in each subband the mixers are pumped differently with possibly opposite balance, the H/V mixer performance is potentially different in each LO subband.

None of the HIFI SIAM entries used during flight made use of this frequency dependence, i.e. corresponding entries in the SIAM were always identical.

Polarisation dependence

Nominally, the H and V mixers are co-aligned. Furthermore the balance of mixer performance is generally good enough to allow for co-adding the spectra observed through the individual mixers channels. Ground measurements however revealed that actual pointing offsets can be as large as 10–20% of the FWHM in bands 1 to 5. In the worst-case scenario of observing point-like sources, one is therefore generally much better off when pointing at an average sky position than at one of the individual mixers and loosing aperture efficiency in the other (which can be as large as 7%). By default we therefore define a synthesized pointing entry, labelled as "S", for co-added H and V spectra for which we point half way between the H and V sky positions. Apart from the default S entry in the SIAM we use the H mixer as the second sub-entry. This choice is based on the fact that the H mixer is on average the most sensitive mixer channel and because the comb-generator of the WBS-H spectrometer was strongest and was expected to have the longer lifetime. We refer to Section 5.8.4 for a detailed discussion of what differences one might observe in the two polarisations. See also Figure 5.44 for an illustration of the actual misalignments between H and V apertures in the various mixer bands.

Chopper dependence

For observing modes not employing focal-plane chopper functionality the default position is the central position on M3 (pick-up mirror) in Fig 4.2. The (primary) chopped position (the position from where the beam is thrown by $1.5'$ at the sky to the other (secondary) chopped position) is chosen to coincide with the mechanical rest position of the focal plane chopper. This mechanical rest position of the chopper nominally coincides with the "-Z" or $+2.4^\circ$ position (towards the internal Calibration Source Assembly) in the HIFI focal plane.

According to the SIAM ICD the HIFI SIAM entries are labelled by Hnm_s where n, m and s are indices of the matrix. Note that s refers to a Solar Aspect Angle (SAA) dependence of the pointing, this was never exploited during the mission and the value of s is therefore 0 for all existing versions of the HIFI SIAM entries. The main indices into the SIAM matrix used by HIFI are therefore n and m. We use n to refer to the active mixer band. The remaining index m (non-zero positive integral index) is composed on the basis of the bit pattern of a 3-bit number. If we denote the least-significant bit by p, the next significant bit by f and the most-significant bit by c, then $m = 1 + c \times 2^2 + f \times 2^1 + p \times 2^0$, where p = polarisation (0 = S, 1 = H), f = frequency (0 = LO subband A, 1 = LO subband B) and finally c = chopper (0 = M3 center, 1 = chopped).

The HIFI aperture naming schema can therefore be summarised as follows:

- aperture coding Hnm_s
- default $s = 0$ (no SAA dependence)
- index n = HIFI band (1, 2, ..., 7)
- index m = 1 + 3 bit number $= 1 + c \times 2^2 + f \times 2^1 + p \times 2^0$ with:
  - least-significant bit p = polarisation (0 = S, 1 = H)
  - next significant bit f = frequency (0 = LO subband A, 1 = LO subband B)
  - most-significant bit c = chopper (0 = M3 center, 1 = chopped)

Table A.1: HIFI aperture identifier description along index m. The index n runs from 1–7, denoting the HIFI mixer band.

Aperture identifier	m	Bit pattern	Chopper	Frequency	Polarization
Hn1_0	1	000	M3 center	LO subband A	S(ynthesized)
Hn2_0	2	001	M3 center	LO subband A	H(orizontal)
Hn3_0	3	010	M3 center	LO subband B	S(ynthesized)
Hn4_0	4	011	M3 center	LO subband B	H(orizontal)
Hn5_0	5	100	Chopped	LO subband A	S(ynthesized)
Hn6_0	6	101	Chopped	LO subband A	H(orizontal)
Hn7_0	7	110	Chopped	LO subband B	S(ynthesized)
Hn9_0	8	111	Chopped	LO subband B	H(orizontal)

The HIFI aperture identifiers Hnm_s have descriptions as listed in Table A.1 assuming no SAA dependence exists.

As an example, this is how an entry in the SIAM for the aperture corresponding to M3 center, LO subband A, S(ynthesized) polarisation in band 1 looked like after a change on April 12, 2010:

```
H11_0 2010-04-12T13:00:33Z
+9.9999587276753230D-01 +2.8604513329032874D-03 -2.6882349837334010D-04
-2.8604512200403084D-03 +9.9999590890071690D-01 +8.0431856198463060D-07
+2.6882469930382137D-04 -3.5358738477422526D-08 +9.9999996386663930D-01
```

Appendix B

Corrupted Data-Frames

On very rare occasions, anomalous spectra can end up in the Level 0 data of an observation context. These corrupted data are normally due to a malfunction in the On-Board Software at the time of transferring data from the spectrometer into the On-board Memory, leading to the loss of parts or all of the affected data. These issues are usually limited to only one or two spectrometers in a given observation.

We have identified the following categories of anomalous data-frames:

1. *Ghost Data-Frames*: these are spectra that exist in a duplicated form in the Level 0 data. The first incarnation of this duplication is usually incomplete, while the second one is complete and can be safely used.
2. *Corrupted Data-Frames*: these are unique spectra that miss part of the information, typically in one or more of the spectrometer subbands (channels filled with zero intensity). In the case of the HRS, this leads to Auto-Correlation spectra partially filled with zeroes, than can therefore not be properly Fourier Transformed into a Level 0.5 spectrum. These data are lost and cannot be recovered
3. *Corrupted Local Oscillator Housekeeping (LO HK)*: these are unique spectra that contain valid science data but hold an erroneous LO readout, usually a frequency that does not make sense for the band in use, or is even outside of the HIFI operational range. These data can be recovered.
4. *Aborted Telecommands*: in the early phases of the mission, inaccuracies in the timing of the uplinked telecommands (TC) could lead to sporadic TC abortion. When this applied to the TC in charge of acquiring science data, those latter were never completed and are thus missing from the Level 0 products. Missing data can also occur during Single Event Upsets (SEU) and are discussed in the next sub-section. These data are lost and cannot be recovered.

At this stage, we have identified $\simeq 50$ Obsids concerned by one of the above anomalies. As a recovery measure, the HIFI pipeline has been adapted in order to deal with those corrupted data in such a way that they get ignored in processing steps beyond the Level 0. Data-Frames falling into category 1 will simply be erased from the Level 0, while for Data-Frames in category 3, the correct LO entries are overhauled in the data by the pipeline based on a look-up table. Data-Frames in category 2 will be stamped with a dedicated row flag `BAD_DATA`. Finally for Data-Frames in category 4, a mitigation

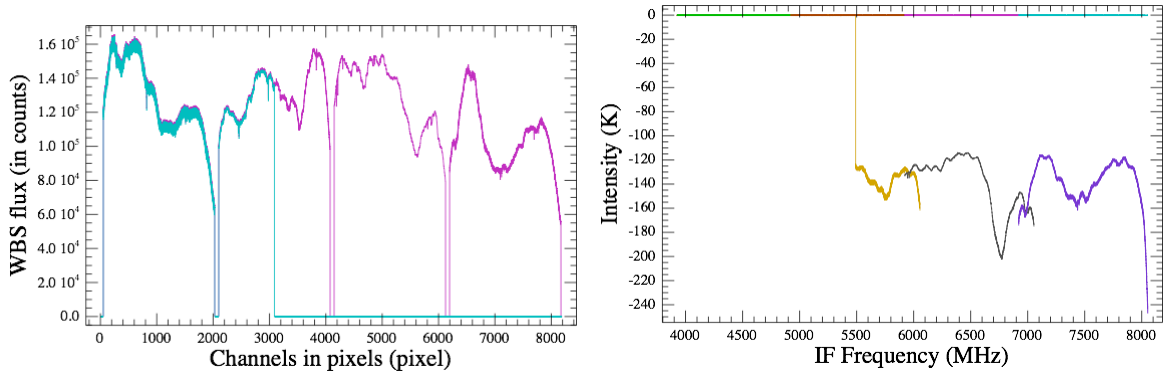

Figure B.1: Example of a corrupted Data Frame. The left panel illustrates total power spectra at Level 0 where part of the channels are set to zero (cyan spectrum) compared to the uncorrupted data in magenta. The right panel shows how the data corruption will manifest in the Level 1 data.

measure consists in flagging those valid data that should have been combined to the missing data. In all cases, this flag will cascade down to the Level 1 or 2 data in any spectrum combination involving the corrupted Data-Frame(s), and the following will then happen in the final products:

- for Point Mode observation, the flagged data at Level 1 will be ignored in the final averaging to create the Level 2 spectrum
- for Spectral Scan observation, the flagged data will be present in the Level 2 but the deconvolution algorithm will ignore them to build the Single Sideband solution at Level 2.5
- for Mapping observation, the flagged data will be present in the Level 2 but the gridding algorithm will ignore them to build the cubes at Level 2.5

The net effect is null only for those of the first and third category since no valid data is lost. For the other categories, the flagging of the corrupted Data-Frames will propagate to the data combined with it in later pipeline stages (e.g. complementary chopping or frequency-switched phases) and will imply a noise degradation in the final Level 2 and 2.5 data.

Table B.1 lists all affected Obsids recognised to date. The detailed table of corrupted Data-Frames is stored in the Calibration Tree (*CleanDF* table – see Section 6.5).

Obsid	Affected backend	Data corruption category
1342210032	WBS-H	Ghost science data-frame
1342216806	WBS-H	Ghost science data-frame
1342219255	WBS-H	Ghost science data-frame
1342219454	WBS-H	Ghost science data-frame
1342230216	WBS-V	Ghost science data-frame
1342201597	HRS-V	Corrupted science data-frame
1342205542	HRS-H	Corrupted science data-frame
1342232737	HRS-V	Corrupted science data-frame
1342235065	HRS-H	Corrupted science data-frame
1342210779	WBS-V	Corrupted science data-frame
1342216804	WBS-V	Corrupted science data-frame
1342202026	WBS-H	Corrupted science data-frame
1342233988	WBS-V	Corrupted science data-frame
1342216320	WBS-V	Corrupted science data-frame
1342216351	WBS-H	Corrupted science data-frame
1342235802	WBS-H	Flag two noisy science data-frames
1342235800	WBS-H	Flag bad Hot-Cold data-frames
1342220517	WBS-H, WBS-V	Flag bad Hot-Cold data-frames
1342201783	WBS-H	Flag bad Hot-Cold data-frames
1342189621	WBS-H, WBS-V	Flag isolated chopper frame
1342210764	WBS-V	Corrupted science data-frame
1342203745	WBS-V	Corrupted science data-frames and TC failures
1342245345	WBS-H, WBS-V	Corrupted science data-frames
1342196567	WBS-H	Flag bad Hot-Cold data-frames
1342239625	WBS-H	Flag bad Hot-Cold data-frames
1342245299	HRS-H	Ghost science data-frames
1342246494	HRS-H	Corrupted science data-frames
1342246544	WBS-H	Corrupted science data-frames
1342246507	WBS-V	Ghost science data-frames
1342245997	WBS-H	Ghost data-frame
1342247204	HRS-H	Corrupted data-frame
1342248767	WBS-H	Corrupted data-frame
1342250406	WBS-V	Ghost data-frame
1342250459	HRS-V	Corrupted data-frame
1342250721	WBS-V	Ghost data-frame
1342250217	WBS-H	Corrupted data-frame
1342250605	HRS-H	Corrupted data-frame
1342257859	WBS-V	Corrupted data-frame
1342258206	WBS-H, WBS-V	Corrupted data-frame
1342256427	WBS-H	Corrupted data-frame
1342265972	WBS-H	Ghost data-frame
1342263227	WBS-V	Ghost data-frame
1342268776	HRS-V	Corrupted data-frame
1342233897	HRS-V	Corrupted LO
1342233898	HRS-H, HRS-V	Corrupted LO
1342190764	WBS-H	Corrupted LO
1342194735	WBS-H	Corrupted LO
1342180555	All backends	TC failures
1342181163	WBS-H, WBS-V	TC failures
1342192563	All backends	TC failures

Table B.1: Compilation of Obsids where corrupted HIFI data have been identified.

Appendix C

Single Event Upsets

Occasionally, Cosmic Ray hits in the memory boards of the various Herschel Control Units could lead to bit flips and trigger a temporary software malfunction. These events are known as *Single Event Upsets* (SEUs). SEUs have turned out to be more frequent than anticipated in the space environment of the Lagrange 2 point, and in the case of the HIFI instrument, they even lead to the loss of the Prime side of the instrument when the Local Oscillator Control Unit (LCU) got hit by a charged particle on August 3, 2009 (OD-81, see Jellema et al. 2010). Since then, numerous SEUs have been recorded in all Herschel instruments, most of them hitting the HIFI LCU. In order to protect the HIFI Redundant side from suffering the same sort of hardware loss as OD-81, the On-Board Software was updated in order to put the instrument in a safe mode whenever a bit flip was detected. When this occurred, the corresponding observations scheduled after the detected event would be lost (although they were usually re-scheduled at a later stage of mission if visibilities allowed it).

In total, we have recorded 71 occurrences of an SEU of some sort in the HIFI instrument. Observations affected by an SEU are stamped as ”FAILED” in the archive. Note that when an Operational Day is affected by an SEU, it does not necessarily imply that all observations from that day have been lost. More details about those events and the particular Operational Days they affected can be found in the Herschel Mission & Satellite Overview Handbook.