A unified Allan variance computation scheme

V. Ossenkopf

SRON National Institute for Space Research,
P.O. Box 800, 9700 AV Groningen,
Netherlands

July 31, 2003

Abstract

We propose a new approach for the computation of the Allan variance of spectrometer data combining the advantages of the two existing methods into a unified scheme. It allows to analyse the stability of an instrument with respect to total-power and spectroscopic fluctuations within the same framework. The method includes an explicit error estimate both for the individual Allan variance spectra and for the derived stability time.

A new definition of the instrument stability time allows to characterise the instrument even in the case of a fluctuation spectrum shallower than $1/f$, as measured for the total power fluctuations in high-electron-mobility transistors. From the stability time and the spectral index of the drift contributions measured by the Allan variance spectrum, optimum cycle lengths for switching between source and reference observations are derived, resulting in a minimum total uncertainty of the resulting data due to radiometric and drift noise.

A first analysis of test measurements for HIFI indicates that the astronomical observations will face huge differences between the stability times relevant for measurements aiming at an accurate determination of the continuum level and for purely spectroscopic measurements.

Contents

1 Introduction **1**

2 Data handling **2**
2.1 Normalisation across spectrometers 2
2.2 Spectroscopic versus total power normalisation 2

3 Computation of the Allan variance **5**
3.1 Convolution schemes 5
3.2 Error estimate 8
3.3 How to characterise the whole spectrometer? 9

4 Interpretation of Allan variance spectra **12**
4.1 Comparison to radiometric noise 12
4.2 Definition of the stability time 14

5 Optimisation of the observations **15**

6 Conclusions **19**
6.1 The optimum method 19
6.2 The expected instrument stability for HIFI 20

References **20**

1 Introduction

The Allan variance (Allan 1966) is a useful measure for the stability of general radio-astronomical equipment, in particular for systems consisting of heterodyne receivers and spectrometer backends (e.g. Kooi et al. 2000). For HIFI the determination of the Allan variance of the instrument is one of the end user requirements for the interactive analysis (Roelfsema et al. 2002). The Allan variance spectrum can be computed from any sufficiently long time series of spectrometer dumps during which all instrumental parameters should be kept constant. The integration time used for a single dump has to be small enough to resolve all possible instrumental instabilities.

The algorithm for the computation of the Allan variance spectrum has to fulfil four requirements:

- Detection and characterisation of all instabilities and of their spectral variation across the measured frequency range
- Efficient use of the measured data to extract a maximum of information from a limited time series
- Providing a measure for the uncertainty of the analysis itself, i.e. an inherent error estimate
- Possibility of a fast implementation to allow the use of the Allan variance analysis as part of a quick-look analysis of measured data

At present, the document on the end user requirements for the HIFI interactive analysis proposes two different algorithms to compute the Allan variance. They use a different mathematical representation of the measured data and fulfil the requirements stated above only partially. The spectroscopic Allan variance as proposed by Schieder et al. (1985) uses only one or two arbitrarily selected channels, thus neglecting a large amount of the measured information and the baseline Allan variance as proposed by Siebertz (1998) allows no identification of problematic channel ranges within the backends and no analysis of total-power drifts.

Here, we propose a new scheme which unifies the different approaches into a single mathematical description and fulfils the four requirements to a large extent. The request for a complete characterisation of the spectral behaviour of the possible instabilities will be translated here into the need for a computation of the Allan variance independently for each backend channel, so that channel by channel variations can be detected, the influence of standing wave instabilities becomes visible, and regions of instabilities across the IF band can be identified.

The request for an efficient data use is fulfilled by actually using the data from all spectrometer channels and by taking into account all possible statistically independent samplings of the temporal behaviour in the analysis. This efficient data use helps to shorten the actual time needed for the measurement of the data. Although a complete characterisation of the instrumental stability requires very long time series of measurements, the derivation of the Allan variance minimum and the drift index of fluctuations at time scales in the order of the Allan minimum time can be obtained already from a measurement that lasts only about three Allan minimum times. Focusing on these two quantities being the actually limiting factors for the planning and the calibration of

astronomical observations allows to draw significant conclusions also from reasonably short time series, thus saving observing time.

Even if the time series of measurements is too short to guarantee a complete statistical invariance of the data, we can derive an explicit error estimate for the Allan variance from the counting statistics of the data taking. With the new definition of the Allan time proposed here this also provides a direct measure for the error of the Allan stability time. The request for a fast implementation is fulfilled by the proposed convolution schemes for the measured data, either in the ordinary time domain or in the Fourier domain.

2 Data handling

2.1 Normalisation across spectrometers

To perform an instrumental stability analysis a time series of spectral dumps has to be taken consisting of spectrometer count rates $c_i(t_k)$, where the index i denotes the channel number across the backend and t_k gives the time for the spectral dump with index k. Here we assume that the actual readout time to perform the dump is negligible relative to the integration time between them. For the HIFI backends this is guaranteed in most observing modes. Only for fast-chop observations using the WBS, modifications to the procedure presented here are required. The integration time between two spectral dumps should be small compared to all instrumental instabilities.

The actually measured spectrum of counts per channel c_i on any source is mainly determined by the bandpass of the instrument which is not flat but strongly varying across the spectrometer bands. All measured astronomical data are thus normalised by this bandpass to deduce the actual input signal. Hence, the stability analysis should use the same kind of normalisation instead of working on the raw backend count rates.

However, a full astronomical calibration including two thermal calibration sources is neither practical nor necessary for the stability analysis. Here, the normalisation scale can be simply defined by the signal level and the zero level of the instrument. A useful normalisation of the spectra is thus provided by

$$s_i(t_k) = \frac{c_i(t_k) - z_i}{\langle c_i(t_k) - z_i \rangle_k} \tag{1}$$

where z_i is the zero level of channel i and the temporal average of each channel is used to normalise the signal level of that channel. This corresponds to the normalisation used in the baseline Allan variance analysis by Siebertz (1998). It provides an approximate equivalence of all backend channels so that differences in their mutual behaviour appear on the same scale. It also has the advantage that all variations are measured relative to the signal level so that they can be compared directly to calibration errors and the radiometric noise level.

2.2 Spectroscopic versus total power normalisation

Most variations in the amplifiers or other components of the signal path lead to fluctuations which are constant across the whole bandpass. Differencing observations which do

not switch between source and reference on a time much shorter than the corresponding fluctuations may then see constant baseline offsets in the calibrated data. However, most astronomical heterodyne observations are eventually not intended for an accurate determination of the continuum level but for the measurement of lines on top of a constant baseline which is typically taken to be zero. In this case only fluctuations which do not influence all channels in the same way result in a degradation of the calibrated astronomical data, often seen as ripples or steps in the baseline.

The astronomical observations thus call for two different kinds of instrument stabilities, the total power stability and the spectroscopic stability. This has to translate into two different types of Allan variance analyses. The total power Allan variance has to trace all instrumental variations as seen by the backends. It is thus computed directly from the normalised spectra as given in Eq. (1).

The spectroscopic Allan variance should only measure instabilities deviating from a common gain variation across the whole band. In the original definition of the spectroscopic Allan variance by Schieder et al. (1985) this is accomplished by considering the difference signal between two channels i and j. This approach, however, suffers from the arbitrariness of the selection of these channels and the impossibility to distinguish between the variations in both channels. To avoid these problems, we propose to use the average across the backend to subtract the continuum level fluctuations. Then, we have to extend Eq. (1) by this difference for the spectroscopic Allan variance computation:

$$s_i(t_k) = \frac{c_i(t_k) - z_i}{\langle c_i(t_k) - z_i \rangle_k} - \left\langle \frac{c_i(t_k) - z_i}{\langle c_i(t_k) - z_i \rangle_k} \right\rangle_i \tag{2}$$

This spectroscopic normalisation corresponds to the subtraction of a zeroth order baseline in the Allan variance method proposed by Siebertz (1998). For long time series, the resulting Allan variance spectra are mathematically also equivalent to the average of two-channel Allan variance spectra taken over all reference channels j, except for the contribution from the considered channel i to the average which is not used in the two-channel spectroscopic Allan variance.

Fig. 1 demonstrates the effect of the spectroscopic normalisation by plotting a time series of normalised spectrometer data $s_i(t_k)$ both for the total power and the spectroscopic approach. It is clearly visible that most variations occur equally in all channels so that they are mainly visible in the total power data. However, a detailed inspection reveals also some differential variations resulting in weak structures visible in the spectroscopically normalised plot.

Because the different subbands of the spectrometers see different signal chains they may show a different stability behaviour. Thus all averages over the channels i considered here, were computed only within the spectrometer subbands. Only if all instabilities in the backends can be completely neglected a treatment of the spectrometer as a single backend is possible.

Altogether we obtain two different kinds of Allan variance plots from each measurement. When performing the analysis at the normalised data given by Eq. (1) we measure the total power stability. This result has to be used when setting up observations which aim for a determination of the continuum level like measurements of absorption lines. With the spectroscopically normalised data from Eq. (2) we measure the spectroscopic stability, which will considerably exceed the total power stability. These results can be used for observations which do not aim for an accurate determination of the

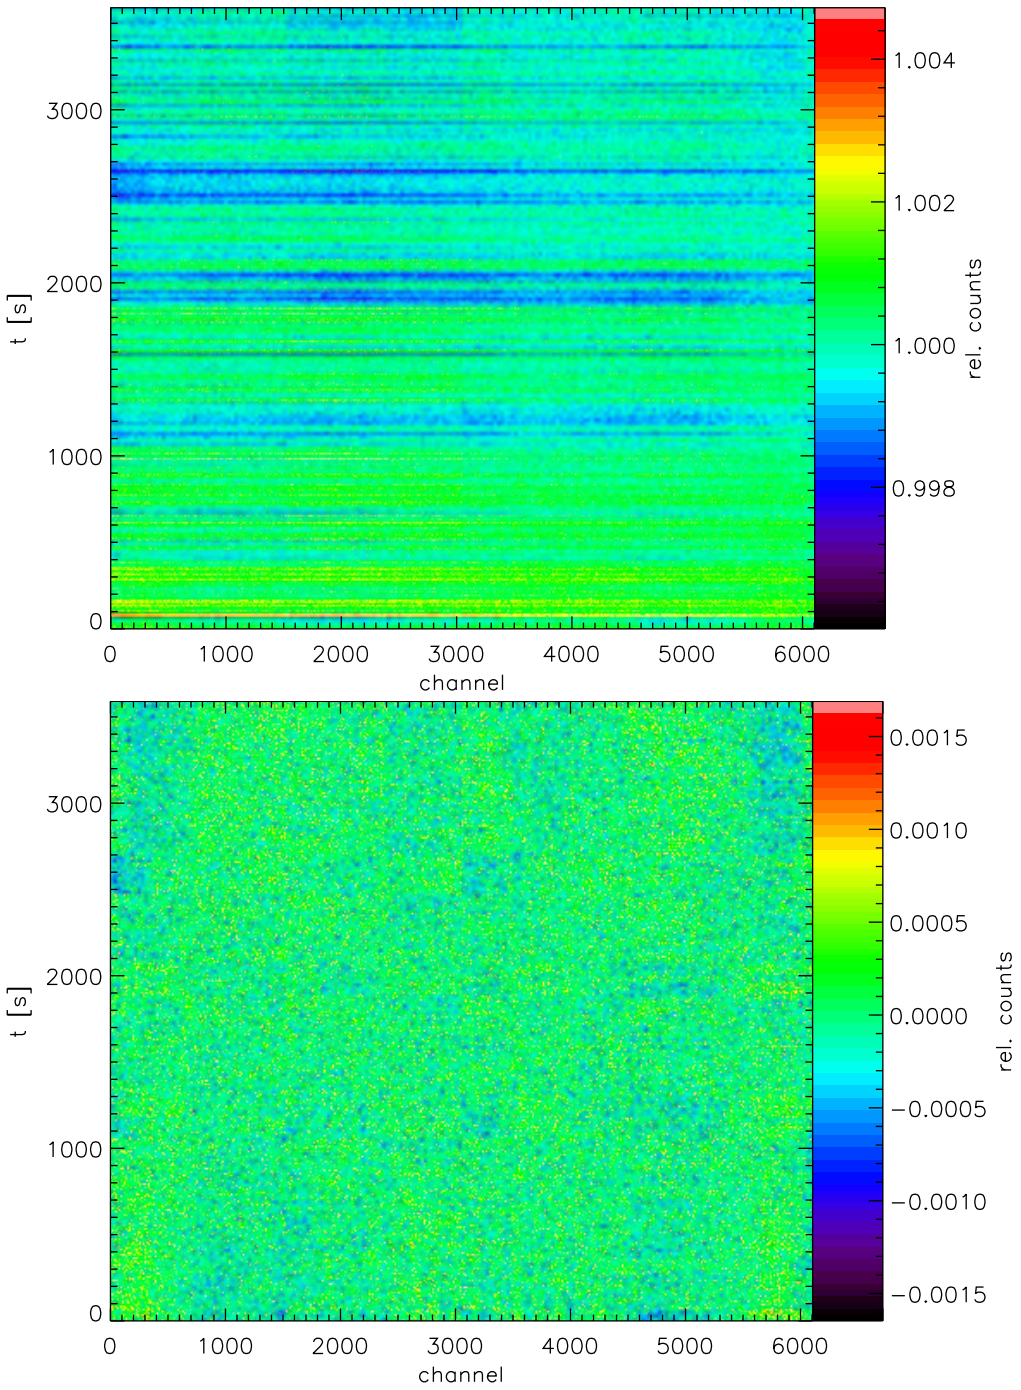

Figure 1: Time series of spectrometer data where either the total-power normalisation (Eq. 1, upper plot) or the spectroscopic normalisation (Eq. 2, lower plot) is applied. The data were taken in stability measurements of an experimental setup using the high electron mobility transistors (HEMTs) foreseen for the HIFI band 2 mixer together with the wideband spectrometer (WBS) prototype in a laboratory measurement in Cologne in May 2002. The four WBS subbands are plotted like one big spectrometer where channels 1-1600 belong to the first subband, 1601-3200 to the second subband, 3201-4800 to the third subband, and 4801-6400 to the fourth subband. Due to the finite resolution of printing only every tenth spectrometer channel is actually displayed.

continuum level, i.e. observations of molecular emission lines on a negligible continuum background. Because the second case leads to much more efficient observations, the selection of the corresponding mode has to find a balance between the need for an accurate continuum level and the request for a high observing efficiency.

3 Computation of the Allan variance

3.1 Convolution schemes

The computation of the Allan variance consists in principle of a convolution of the signal data $s_i(t_k)$ from Eqs. (1) or (2) by a Haar wavelet

$$\sqcap\!\sqcup_L = \begin{cases} 1/L & \text{for } -L \le t < 0 \\ -1/L & \text{for } 0 \le t < L \\ 0 & \text{everywhere else} \end{cases} \tag{3}$$

of size L and the computation of the variance of the convolved signal[1, 2]:

$$\sigma^2_{\mathrm{A},i}(L) = \left\langle \left(s_i(t_k) * \sqcap\!\sqcup_L - \left\langle s_i(t_k) * \sqcap\!\sqcup_L \right\rangle_k \right)^2 \right\rangle_k \tag{4}$$

Plotting the Allan variance $\sigma^2_{\mathrm{A}}(L)$ as a function of the filter size L shows the variation of the signal on the scale of the temporal lag L. The filter lengths L should be chosen to be integer multiples of the step size $\Delta t = t_{k+1} - t_k$, i.e. $L = l \times \Delta t$ with $l = 1 \ldots l_{\max}$. The maximum filter size still providing reasonably reliable Allan variance data falls between about one sixth and one third of the total length of the time series.

Computing the Allan variance by actually convolving the time series by the $\sqcap\!\sqcup$-filter function for each time step, however, seems to be a waste of computing time, because the neighbouring values in the convolved time series are no longer statistically independent. Only when considering points separated by a lag in the order of the filter size L the result of the convolution becomes independent again. A more efficient method is thus to compute the convolution integral only for filter settings separated by the filter size. This approach was used by Schieder et al. (1985) when using the average over a

[1] Please, note that the original definition of the Allan variance is smaller by a factor 1/2. We omit this factor to allow a direct comparison to the drift error in observations.

[2] The variance definition used here is

$$\sigma^2 = \left\langle (x_k - \langle x_k \rangle_k)^2 \right\rangle_k = \frac{1}{N} \sum_{k=1}^{N} (x_k - \langle x_k \rangle_k)^2$$

This deviates from the ordinary variance definition

$$\sigma^2_{\mathrm{standard}} = \frac{1}{N-1} \sum_{k=1}^{N} (x_k - \langle x_k \rangle_k)^2$$

for small numbers N. It has, however, the advantage that it disregards from the exact way of sampling a given continuous distribution, measuring only the internal properties of the distribution, as long as the sampling is dense enough.

reduced number of points to compute the Allan variance

$$\sigma^2_{\mathrm{A},i}(L) = \left\langle \left(S_i(K) - S_i(K+1) - \langle S_i(K) - S_i(K+1)\rangle_K\right)^2 \right\rangle_K \tag{5}$$

with

$$S_i(K) = \frac{1}{l} \sum_{k=Kl+1}^{(K+1)l} s_i(t_k) \tag{6}$$

Compared to the full convolution (Eq. 4) this corresponds to counting only points separated by L in the convolved function $s_i(t_k) * \sqcap\!\sqcup_L$ when computing the variance. Each data point of the signal contributes twice – once in the positive and once in the negative term for the binned signal.

However, this sum does not use all information contained in the time series on the considered scale, because an alternative binning of the data shifted by half the filter length L is possible providing independent information which is lost when using only the binning from Eq. (6). To exploit the full information of the time series on the scale L, the convolved function $s_i(t_k) * \sqcap\!\sqcup_L$ has to be sampled on a raster of $L/2$, not only L, when computing the Allan variance. The full sampling can be expressed by an additional term in the sum

$$\begin{aligned} \sigma^2_{\mathrm{A},i}(L) \;=\; & \Big\langle \left(S_i(K) + T_i(K) - S_i(K+1) - T_i(K+1)\right. \\ & \left. - \left\langle S_i(K) + T_i(K) - S_i(K+1) - T_i(K+1)\right\rangle_K\right)^2 \Big\rangle_K \end{aligned} \tag{7}$$

with

$$T_i(K) = \frac{1}{l} \sum_{k=(K+1/2)l+1}^{(K+3/2)l} s_i(t_k) \tag{8}$$

Fig. 2 demonstrates the effect of the sampling by comparing the Allan variance spectrum for a time series of an arbitrary spectrometer channel computed from the full convolution of the time series with the spectra obtained from the two different samplings discussed above. It is interesting to note that the results of the fully sampled computation still show a considerable scatter around the smooth curve provided by the convolution at each point. Especially at the largest lags – visible here above 50 s – we find always noticeable deviations resulting from the fact, that the discrete sampling approach ignores a relatively large fraction of the data when the time series does do not match multiple integers of the filter size and the filter itself is large. Nevertheless, almost all values fall within the statistical error bars shown for the full convolution and computed below. The results from the half-sampled computation scatter much more requiring long time series for a reliable determination of the Allan-variance spectrum.

We can conclude that the brute-force approach of the full convolution of the time series with the filter function provides the best results by not neglecting any data and guaranteeing that every feature in the spectrum is covered by an appropriate filter setting. This results in smooth curves of the Allan variance spectra facilitating a fit and interpretation even by the naked eye. Using today's computer technology such a convolution is possible by means of a Fast-Fourier transform for spectrometer time series of up to a few thousand steps. For longer time series, the numerically simpler approach

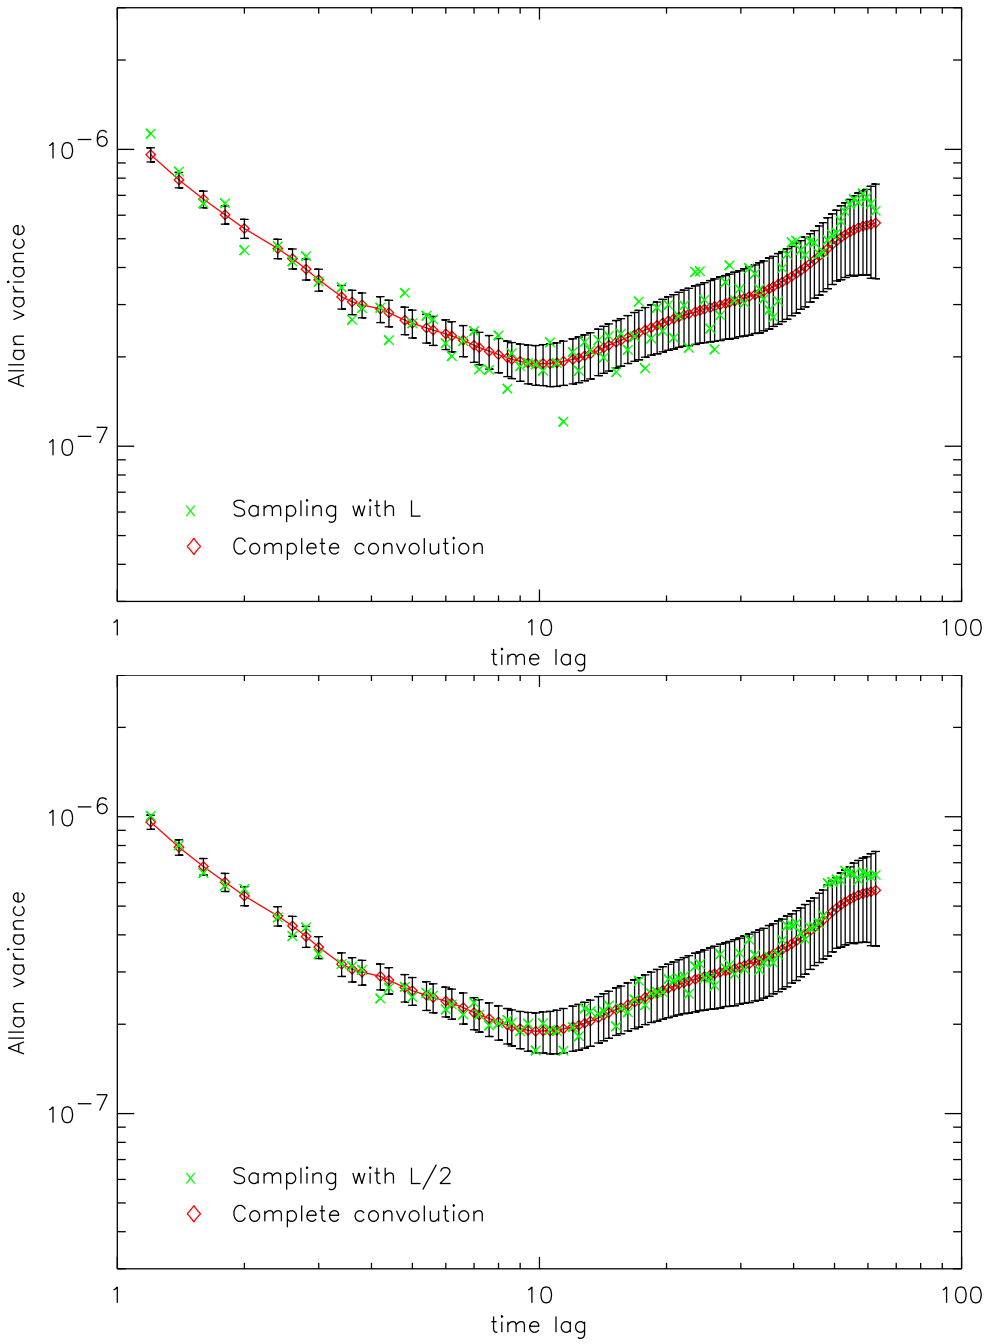

Figure 2: Allan variance spectrum of a single channel computed by the full convolution of each point of the time series with the ⎍-wavelet (diamonds) compared with the spectra obtained from a half sampling of the convolved time series (upper plot) and from the full sampling (lower plot). The error bars shown for the results from the full convolution integral are determined by the statistical uncertainty due to the finite number of data points (see Sect. 3.2).

of the discrete sampling has to be used, however. The full sampling by one point from the convolved time series every $l/2$ steps still provides reliable data completely within the statistical uncertainty inherent to the Allan variance analysis. The sparse sampling of the data as used previously, however, should not be used any more because it wastes

measured data by ignoring part of the contained information and the full sampling is not more complex to implement and compute.

3.2 Error estimate

The previous section has shown that the different possibilities to sample the distribution of data values in the filter-convolved time series lead to a considerable scatter of the resulting Allan variance values. This can be quantified as the statistical error of the Allan variance. We concentrate here on this intrinsic uncertainty of the method neglecting the error propagation of possible uncertainties of the measured data into the Allan variance values.

The statistical error results from the sampling error when scanning a continuous distribution $f(t)$ by taking N data values at discrete randomly selected points K. It is well known that the uncertainty in the determination of the average value of the sampled distribution is given by a Poisson counting error and the variance of the distribution

$$\delta\langle f\rangle_K = \sqrt{\frac{\langle(f-\langle f\rangle_t)^2\rangle_t}{N}} \tag{9}$$

Equivalently, one can derive the uncertainty of the measured variance of the distribution caused by the discrete sampling as

$$\begin{aligned}\delta\langle(f-\langle f\rangle_K)^2\rangle_K &= \sqrt{\frac{\langle(f-\langle f\rangle_t)^4\rangle_t - \langle(f-\langle f\rangle_t)^2\rangle_t^2}{N}} \\ &= \langle(f-\langle f\rangle_t)^2\rangle_t\sqrt{\frac{Kur-1}{N}}\end{aligned} \tag{10}$$

where Kur denotes the kurtosis of the distribution, characterising its fourth moment. Gaussian distributions exhibit a kurtosis value of 3. Exponential distributions show $Kur = 6$. Thus the relative accuracy of the measured variance of a distribution depends mainly on the number of points used to sample the distribution. The kurtosis measuring the relative strength of the wings of the distribution weighs this counting error by a factor of a few.

It is obvious that these principles apply as well to the determination of the Allan variance. As discussed above each distribution of filter-convolved data values can only be sampled in steps of at least $l/2$ to obtain statistically independent values. If the data represent a random series the even sampling in time steps corresponds to a random sampling of the distribution of data values so that the equation above can be applied. This assumption is not always fulfilled in the measurement of drift processes but it is in general justified when sufficiently long time series are measured. Thus we can estimate the statistical error of the Allan variance for a time series with the length $N \times L$ by

$$\delta\sigma_{\mathrm{A}}^2(L) = \sqrt{\frac{\langle(s_i(t_k) * \sqcap\!\sqcup_L - \langle s_i(t_k) * \sqcap\!\sqcup_L\rangle_K)^4\rangle_K - \langle(s_i(t_k) * \sqcap\!\sqcup_L - \langle s_i(t_k) * \sqcap\!\sqcup_L\rangle_K)^2\rangle_K^2}{N}} \tag{11}$$

These error bars were plotted in Fig. 2. The comparison with the variations from the different sampling discussed in Sect. 2 shows that the error bars provide a reasonable feeling for the possible uncertainty from the sampling of the convolved data.

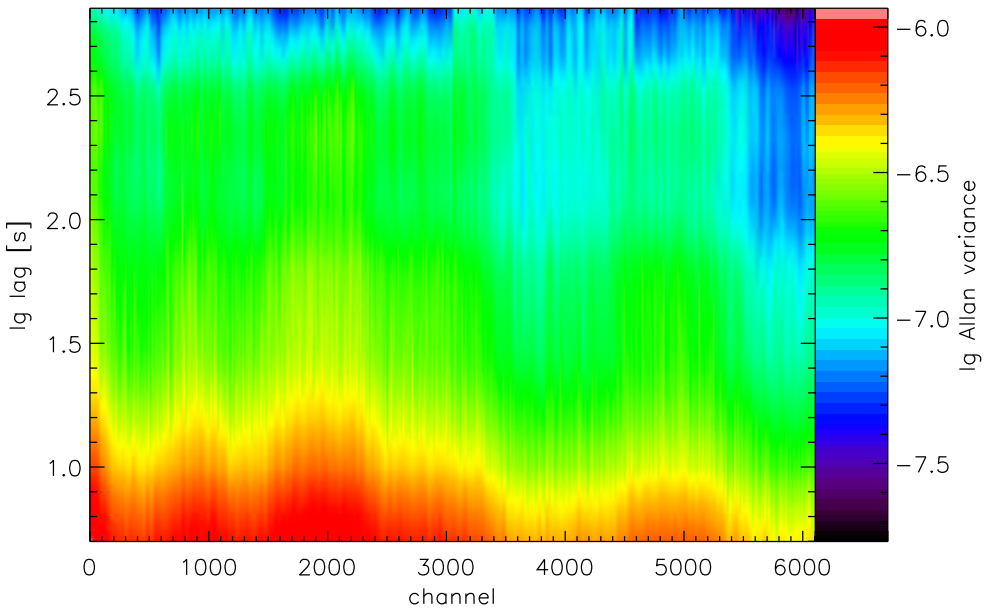

Figure 3: Channel-by channel total-power Allan variance spectra for the test measurement described in Fig. 1. The colour coding shows the logarithm of the Allan variance.

3.3 How to characterise the whole spectrometer?

From the Allan variance of the time series of a single channel as displayed in Fig. 2 it is not obvious how to characterise a complete instrument with thousands of spectrometer channels which are partially, but not completely independent. A first approach is the computation of the Allan variance channel by channel and the visualisation of the result in a three-dimensional plot. This has practical disadvantages that a surface plot is always more difficult to interpreted by eye than a two-dimensional plot, that it is not possible to include information about the error bars in the plot, and that the use of the Allan variance measurement for the optimisation of the observing strategy has to rely on a single number, not a multi-dimensional data set. Nevertheless, this kind of plot is essential to get a feeling for the variability of the stability across a spectrometer.

The result of such a channel-by channel analysis is demonstrated in Fig. 3 showing the total-power Allan variance in logarithmic units for the HIFI test measurement shown in Fig. 1. The wideband spectrometer consists of four subbands which are plotted in a single row. From each subband only 1530 channels were used. The boundary between the second and the third subband is visible in the plot by the sudden change of the Allan variance value at large time lags. From this plot it is obvious that variations of the instrumental behaviour across the spectrometer have to be considered. Not all channels follow the same stability pattern but the variations within each spectrometer subband are relatively small. Thus it is reasonable to characterise the behaviour of the instrument by a single Allan variance spectrum for each subband. A plot like Fig. 3 should nevertheless always by computed to check whether the assumption of the approximate equivalence of all channels is justified.

For the derivation of constraints for the observing strategy it seems plausible to characterise the whole instrument by the properties of the worst, i.e. most instable channel. This is the most reliable approach also applicable to observations where the measure-

ment in all channels is equally important, e.g. in frequency surveys of rich emission spectra with hundreds of lines per spectrum. However, it does not take into account that in most observations the observed lines cover only a very small fraction of the whole spectrometer output whereas bad channels are typically concentrated towards the edges of the IF band. Here, an average Allan variance spectrum makes more sense.

One can consider three different ways of averaging. When starting from the channel-by-channel Allan variance analysis an average Allan variance spectrum is given by

$$\sigma_{\mathrm{A}}^2(L) = \left\langle \left\langle \left(s_i(t_k) * \sqcap\!\sqcup_L - \left\langle s_i(t_k) * \sqcap\!\sqcup_L \right\rangle_{t_k} \right)^2 \right\rangle_{t_k} \right\rangle_i \tag{12}$$

This approach corresponds to averaging the Allan variance spectrum obtained from two channels following the method by Schieder et al. (1985) over all pairs of channels.

A second possible approach is actually used by the baseline Allan variance method by Siebertz (1998). Here, the variance within the convolved spectrum for each time step is considered and in a second step the average over all time steps is performed. We can write this as

$$\sigma_{\mathrm{A}}^2(L) = \left\langle \left\langle \left(s_i(t_k) * \sqcap\!\sqcup_L - \left\langle s_i(t_k) * \sqcap\!\sqcup_L \right\rangle_{i} \right)^2 \right\rangle_{i} \right\rangle_{t_k} \tag{13}$$

This kind of averaging is not able to monitor total-power variations because they enter the data $s_i(t_k) * \sqcap\!\sqcup_L$ in the same way as $\left\langle s_i(t_k) * \sqcap\!\sqcup_L \right\rangle_i$ so that they are subtracted and removed from the Allan variance spectrum. It thus is only useful to measure the spectroscopic Allan variance of an instrument.

Alternatively, one can determine the variance relative to the grand average of the normalised and convolved spectra over the whole data field

$$\sigma_{\mathrm{A}}^2(L) = \left\langle \left(s_i(t_k) * \sqcap\!\sqcup_L - \left\langle s_i(t_k) * \sqcap\!\sqcup_L \right\rangle_{t_k,i} \right)^2 \right\rangle_{t_k,i} \tag{14}$$

The result for the spectroscopically normalised time series of the HIFI test measurement shown in Fig. 1 are demonstrated in Fig. 4. At all time lags below 200 s the different averages show identical Allan variance spectra following approximately the behaviour expected for pure radiometric noise. In contrast, the selection of the worst channel at each time lag shows that even at very short time scales a few channels are already influenced by drift contributions with an amplitude of 30% of the radiometric noise. This grows to the full size of the radiometric noise after 40 s. We have to conclude that the stability time of the instrument is very short if one has to guarantee that the drift contribution remains small for each individual channel of the backend.

Comparing the average of the baseline Allan variance with the Allan variance using the grand average in the difference shows very similar spectra which are hardly to distinguish by eye in the figure. In contrast, the average of the channel Allan variance spectra is always smaller at large lags. This means that the variation of the convolved data relative to the average in the corresponding channel is always smaller than the variation relative to the global average. In contrast the variation relative to the average over the spectrum at a given time step is almost identical to the variation relative to the global

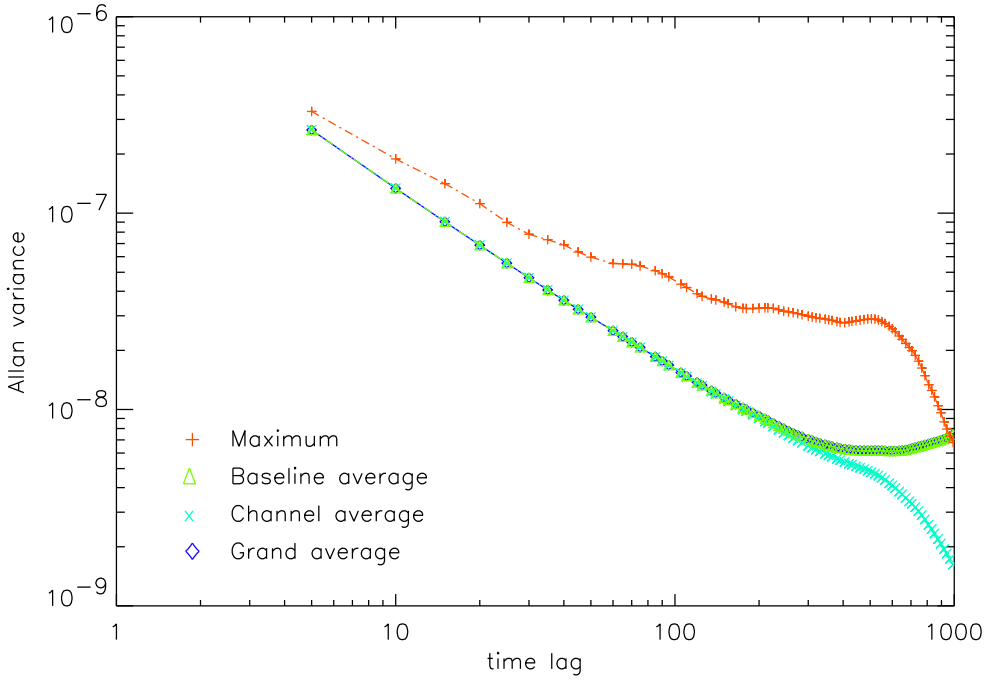

Figure 4: Comparison of the four possible ways to characterise a whole spectrometer band by a single Allan variance spectrum. The spectroscopically normalised data from the first spectrometer band of the tests shown in Fig. 1 have been used. The results when averaging the baseline Allan variance plots and when determining the variance relative to the grand average including simultaneously the channel scale and the time scale are so close that the curves are hardly distinguishable by eye.

average. This can be understood from the the bandpass normalisation and the spectroscopic normalisation of the signal $s_i(t_k)$ discussed in Sect. 2.1. The bandpass normalisation guarantees that average $\langle s_i(t_k)\rangle_{t_k}$ for each channel is identical to the global average $\langle s_i(t_k)\rangle_{t_k,i}$ whereas the spectroscopic normalisation leads to $\langle s_i(t_k)\rangle_i = \langle s_i(t_k)\rangle_{t_k,i}$ for each spectrum. The convolution with the Allan filter now leads to offsets of $\langle s_i(t_k) * \sqcap\!\sqcup_L\rangle_{t_k}$ for the individual channels from the global average when different channels have different trends not cancelling out completely to zero. In contrast, the convolution hardly changes the average of the individual spectra for spectroscopically normalised data, so that $\langle s_i(t_k) * \sqcap\!\sqcup_L\rangle_i \approx \langle s_i(t_k) * \sqcap\!\sqcup_L\rangle_{t_k,i}$ for each time step. For spectroscopically normalised data and sufficiently long time series where the long-term drifts are cancelled out all three averaging methods should provide the same results. However, no measurement will be long enough to actually completely cancel out all drifts, so that one averaging method has to be chosen.

This choice should be based on the practical application of the Allan variance analysis constraining the observing mode timing in such a way that the drift error in measured data can be controlled. All spectroscopic observations are actually interested in the shape of a whole spectrum not regarded as a collection of individual channel data. Drift errors showing up as irregular baseline distortions often result in a clear degradation of the scientific value of the data even if the magnitude of the distortions does not exceed the radiometric noise in the data. Thus the mutual relation of drift contributions across the spectrum has to be taken into account clearly identifying the average

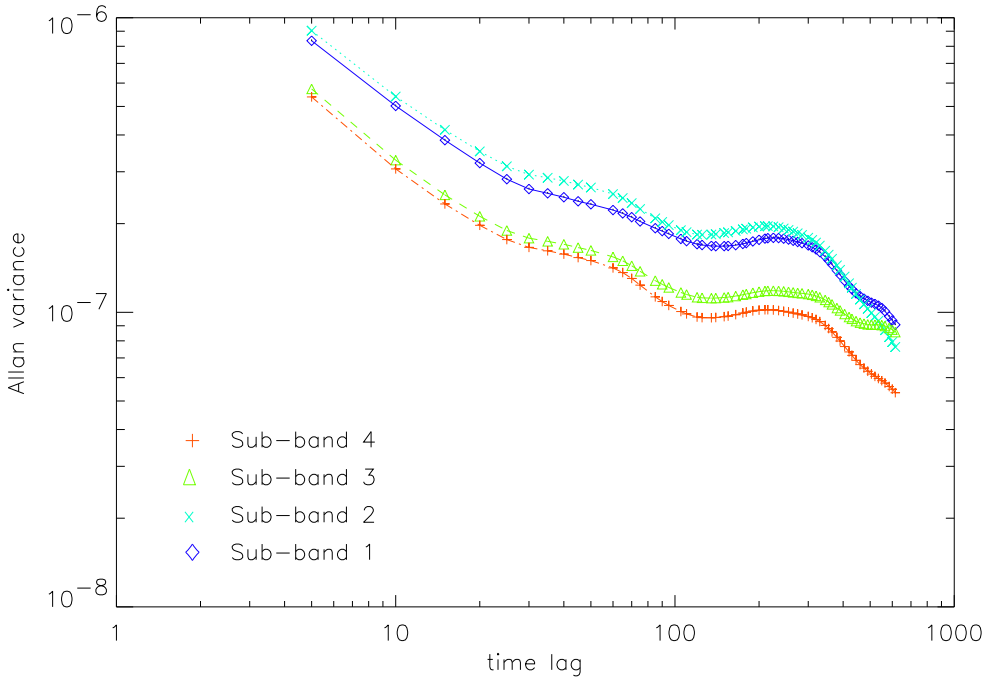

Figure 5: Average total power Allan variance spectra for the different WBS subbands computed for the time series of data presented in Fig. 1.

of the baseline Allan variance or the Allan variance relative to the grand average as the approaches adequate to the needs of astronomical observations. Because the baseline Allan variance is not able to characterise total power drifts we propose to always use the "grand average" method. It combines the advantage of the baseline Allan variance method reflecting the observer's view on spectroscopic data with the possibility to analyse the drift behaviour of an instrument including total-power variations. However, we want to stress again that the characterisation of the instrument by a single Allan variance spectrum is only justified if an inspection of a plot like Fig. 3 has proven that no strong deviations of the drift behaviour across the spectrometer occur.

For setups with an array-spectrometer like the HIFI-WBS each spectrometer subband has to be characterised individually. Fig. 5 shows the resulting total-power Allan variance spectra for the HIFI laboratory measurement already used for demonstration in Fig. 3. One can clearly see that two subbands have a somewhat higher noise over all time scales and that the absolute Allan variance values indicate a much stronger drift contribution than in the spectroscopic case. The total power fluctuations exceed the radiometric noise considerably at all scales. The contribution from spectroscopic fluctuations, which are close to the radiometric noise, as shown in Fig. 4 for the first subband is always negligible.

4 Interpretation of Allan variance spectra

4.1 Comparison to radiometric noise

The Allan variance measurement does not distinguish between fluctuations from the radiometric white noise and drift noise from fluctuations in the instrumental response

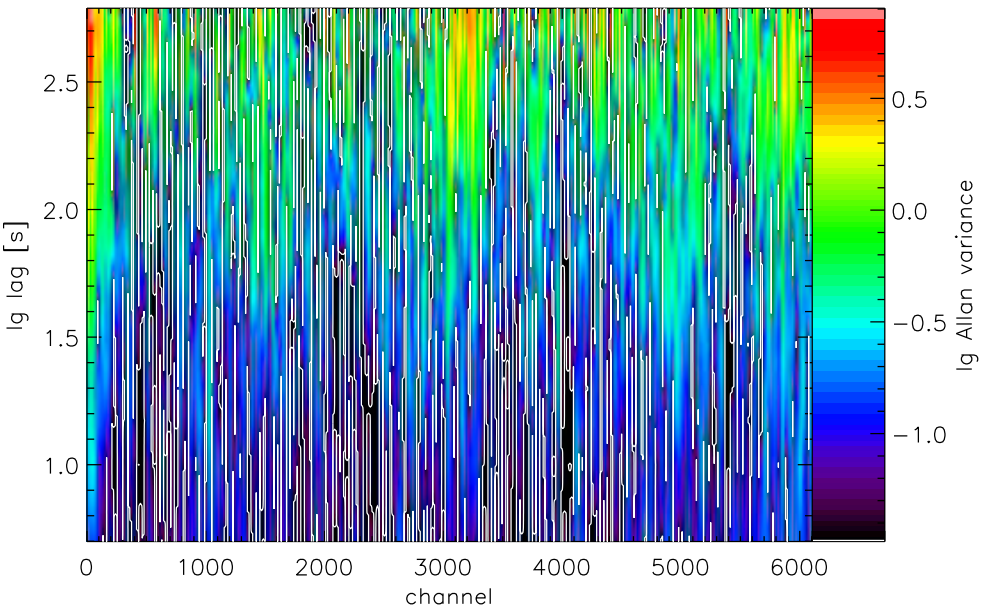

Figure 6: Channel-by-channel spectroscopic Allan variance spectra relative to the radiometric noise level for the specified fluctuation bandwidth of 1.6 MHz. The radiometric noise is subtracted. The white contours indicate parts of the spectrum where the variations drop below the radiometric noise level. The data from Fig. 1 are used so that the figure can be compared to the total power Allan variance plot in Fig. 3.

but contains both contributions. However, they can be separated based on the different spectral characteristics. Fluctuations with a $1/f^{\alpha}$ power spectrum show up in the Allan variance as $L^{\alpha-1}$ spectra (Schieder & Kramer 2001). Assuming white radiometric noise and a power law drift noise with a spectral exponent α gives an Allan variance spectrum

$$\sigma_{\mathrm{A}}^2(L) = \frac{2}{B_{Fl}L} + AL^{\alpha-1} \tag{15}$$

Here B_{Fl} is the fluctuation bandwidth per backend channel determined by the power spectrum of the noise (Kraus 1980) and A characterised the amplitude of the instrumental drift. The radiometric noise contribution can be subtracted as $2/(B_{Fl}L)$ from the Allan variance spectrum to determine the drift contributions.

To derive timing constraints for astronomical observations from the Allan variance spectra the ratio between the drift contribution and the radiometric noise contribution has to be considered because all observations will aim for data where the error due to instrumental drifts is small compared to the radiometric error of the observation. Thus the drift Allan variance normalised to the radiometric noise

$$\sigma_{\mathrm{A,drift}}^2(L) = \frac{B_{Fl}L}{2}\sigma_A^2(L) - 1 = \frac{B_{Fl}A}{2}L^{\alpha} \tag{16}$$

is a useful quantity to describe the actual impact of the drift noise for an astronomical observation.

The result of this noise normalisation is demonstrated in Fig. 6 for the spectroscopic Allan variance of the HIFI test data. We find in the plot some regions where the measured spectrum drops below the radiometric noise limit by a few percent – indicated

here by white contours. These small negative values are not surprising because the fluctuation bandwidth used in the subtraction is an average over the whole spectrometer whereas the actual fluctuation bandwidth varies slightly between individual channels. The average spectroscopic Allan variance spectrum shown in Fig. 4 falls nowhere below the radiometric limit but is very close to it for small lags. Fig. 6 shows, however, also that in some parts of the spectrometer the drift contribution exceeds the noise RMS by a factor 3–8 after 600 s whereas it is smaller than the noise RMS in other parts at the same time.

The situation is different for the total-power Allan variance as shown in Fig. 5. Already at the shortest measured time scale the Allan variance clearly exceeds the radiometric noise contribution of $2.5\,10^{-7}$ there. At lags between 10 and 300 s the spectra of all subbands can be approximated by an $L^{-0.3}$ dependence corresponding to a $1/f^{0.7}$ characteristics of the fluctuations. This is in agreement with stability measurements of the isolated HEMTs by the Centro Astronomico de Yebes (Whyborn 2003), indicating that the total-power stability of the system is mainly determined by their gain fluctuations. In the spectroscopic Allan variance, the time series turns out to be too short for an accurate determination of the spectral index. All α values between about 1.5 and 2.5 are consistent with the measurements.

4.2 Definition of the stability time

It is now possible to define an Allan time as the lag at which the Allan variance spectrums changes from being dominated by radiometric noise to being dominated by the instrumental drift. The traditional definition uses the minimum of the Allan variance spectrum. At smaller lags the fluctuations are dominated by the radiometric noise and decrease with the size of the Allan filter, at larger lags the drift terms dominate resulting in an increase of the fluctuations with filter size. This approach, however, is only applicable if the drift follows a usual spectral characteristics with an index $\alpha > 1$. Otherwise no minimum is formed like (see Fig. 5).

The use of the Allan minimum time t_A thus has two disadvantages: **i)** for drift noise shallower than $1/f$ the Allan variance has no minimum although it is still a good measure for the stability of the system. **ii)** it is difficult to derive an error estimate for the Allan minimum time from the statistical uncertainty of the Allan variance.

Here, we propose another definition of the Allan time t'_A based on the normalisation discussed above: t'_A is the shortest lag where the drift contribution to the per-channel uncertainty equals the radiometric contribution, i.e. $\sigma^2_{\mathrm{A,drift}}(t'_\mathrm{A}) = 1$. This means that at t'_A the total Allan variance amounts to twice the radiometric Allan variance. For a drift noise following a $1/f^2$ spectral dependence, the new Allan time t'_A agrees with the Allan minimum time t_A. The new Allan time definition, however, has two practical disadvantages: **i)** The minimum of the Allan variance can be easily determined even by the naked eye whereas the new definition needs an exact comparison of the drift contribution with the radiometric noise. **ii)** Because most existing measurements characterise the instrument stability by the Allan minimum time, a comparison with their results becomes more difficult – although possible. The translation between both Allan times can be computed from Eq. (15) as

$$t'_\mathrm{A} = (\alpha - 1)^{1/\alpha} t_\mathrm{A} \tag{17}$$

The radiometric noise contribution is always determined by the actual fluctuation bandwidth of a measurement. Binning or smoothing of data to to a lower fluctuation bandwidth is often performed to obtain a lower radiometric noise. Consequently, this has implications for the Allan time of a measurement because the instrumental drift itself is not affected by the smoothing. Resolving Eq. (15) for t'_A gives

$$t'_\mathrm{A} = \left(\frac{2}{AB_\mathrm{Fl}}\right)^{1/\alpha} \tag{18}$$

showing that the Allan time shifts to smaller lags when increasing the fluctuation bandwidth by

$$t'_\mathrm{A}(B_\mathrm{Fl,binned}) = \left(\frac{B_\mathrm{Fl}}{B_\mathrm{Fl,binned}}\right)^{1/\alpha} t'_\mathrm{A}(B_\mathrm{Fl}) \tag{19}$$

This relation is identical to the corresponding relation for the traditional Allan minimum time (Schieder & Kramer 2001).

The uncertainty of the Allan time t'_A can be directly computed from the uncertainty of the Allan variance in the environment of t'_A by

$$\frac{\delta t'_\mathrm{A}}{t'_\mathrm{A}} = \frac{1}{|\alpha|}\frac{\delta\sigma^2_\mathrm{A}(L)}{\sigma^2_\mathrm{A}(L)} \tag{20}$$

This error remains finite even in case of $1/f$ noise whereas it diverges for the traditional Allan minimum time.

For the test measurements with the HIFI HEMTs we obtain for a fluctuation bandwidth of 1.6 MHz total-power Allan times t'_A of a few seconds whereas the spectroscopic Allan times fall between 400 and 900 s.

5 Optimisation of the observations

Schieder & Kramer (2001) used the Allan-variance minimum to derive constraints for the optimum timing of astronomical observations. They considered, however, only drift contributions with a spectral index α between 2 and 3 typical for the drift contribution determined in spectroscopic Allan variance measurements. Using the new knowledge on the Allan variance in total-power fluctuations from Fig. 5 we repeat their computation without this restriction also applying the new Allan time definition discussed above. Corresponding equations for the Allan minimum time were also derived by Schieder (priv. comm.).

A large number of observing modes is characterised by equal time intervals spent on the astronomical source and on the reference. These are position-switch, chopped or frequency switch observations. If they are performed and reduced in an REF-SOURCE-SOURCE-REF scheme, they consist of integrations for a time T on the source, a time T on the reference and a dead time T_d in between. By using the information about the instrumental drift obtained from the Allan variance spectrum, in particular the Allan time t'_A and the spectral index of the drift α it is possible to compute the average drift error in each difference measurement $S_k - R_k$, where S_k is the integrated signal over the

kth source phase and R_k is the integrated signal during the kth reference phase. The total uncertainty of the astronomical measurement is then characterised by the variance

$$\sigma^2_{\rm obs}(T, T_{\rm d}) = \left\langle \left(S_k - R_k - \langle S_k - R_k \rangle_k\right)^2 \right\rangle_k \tag{21}$$

Following the formalism provided by Schieder & Kramer (2001) we can derive this variance normalised to the average signal level from Eq. (15)

$$\frac{\sigma^2_{\rm obs}(T, T_{\rm d})}{\langle s \rangle^2} = \frac{2}{B_{\rm Fl} T} + A \frac{(2T + T_{\rm D})^{\alpha+1} - 2(T + T_{\rm d})^{\alpha+1} + T_{\rm d}^{\alpha+1} - 2T^{\alpha+1}}{2(2^\alpha - 2)T^2} \tag{22}$$

The first term describes the radiometric noise which does not depend on the dead time. The second term is the drift noise. For $T_{\rm d} = 0$ we return to the known relation (15) for the Allan variance with no dead time between subsequent data dumps. Eq. (22) holds for all spectral indices α between 0 and 3 except for $\alpha = 1$ where a logarithmic divergence occurs leading to a somewhat different functional description.

Substituting A by $t'_{\rm A}$ using Eq. (18) and normalising all times relative to the Allan time then gives

$$\frac{\sigma^2_{\rm obs}(x, d)}{\langle s \rangle^2} = \frac{2}{B_{\rm Fl} t'_{\rm A}} \left(\frac{1}{x} + \frac{(2x + d)^{\alpha+1} - 2(x + d)^{\alpha+1} + d^{\alpha+1} - 2x^{\alpha+1}}{2(2^\alpha - 2)x^2} \right) \tag{23}$$

with $x = T/t'_{\rm A}$ and $d = T_{\rm d}/t'_{\rm A}$. The use of the traditional Allan minimum time in this equation would add the factor $1/(\alpha - 1)$ to the drift term. For $\alpha = 2$ and $\alpha = 3$ one can then reproduce the numbers obtained by Schieder & Kramer 2001.

The optimum observing mode is now characterised by a minimum total noise – composed of radiometric and drift noise – obtained in a given observing time. With $t_{\rm tot}/(2T + T_{\rm d})$ source-reference pairs in a total observing time $t_{\rm tot}$ we obtain the total noise of the observation as

$$\frac{\sigma^2_{\rm tot}(x, d)}{\langle s \rangle^2} = \frac{4x + 2d}{B_{\rm Fl} t_{\rm tot}} \left(\frac{1}{x} + \frac{(2x + d)^{\alpha+1} - 2(x + d)^{\alpha+1} + d^{\alpha+1} - 2x^{\alpha+1}}{2(2^\alpha - 2)x^2} \right) \tag{24}$$

(see Eq. 9). The behaviour of this total noise is shown in Fig. 7 for a drift spectral index of 0.7 corresponding to the total-power measurements discussed in Sect. 4.1. The noise RMS relative to the radiometric noise in case of no dead times, i.e. $\sigma_{\rm tot}(x, d)/\langle s \rangle \times \sqrt{B_{\rm Fl} t_{\rm tot}/4}$, is plotted as a function of the integration time per cycle relative to the Allan time $t'_{\rm A}$ for different relative dead times. One can see that the minima defining the optimum integration time are much wider than the corresponding minima computed by Schieder & Kramer (2001) for a drift spectral index of 2. But the optimisation of the observing mode based on the proposed Allan time is easily possible, even in this case, where the Allan variance spectrum itself does not show a minimum.

Consequently, we can compute the optimum integration time per cycle relative to the Allan time as a function of the dead time and the spectral index of the drift contributions by

$$\begin{aligned} x_{\rm opt} = t_{{\rm A},\Delta\nu} \times {\rm root}_x \big\{ & (2x + d)^{\alpha+1}(\alpha x - d) - 2(x + d)^{\alpha+1}(2x + d)(\alpha x + d) + (x + d)^{\alpha+1} d \\ & - x^{\alpha+1}[\alpha(2x + d) - d] - d^{\alpha+1}(x + d) - (2^\alpha - 2) d x \big\} \end{aligned} \tag{25}$$

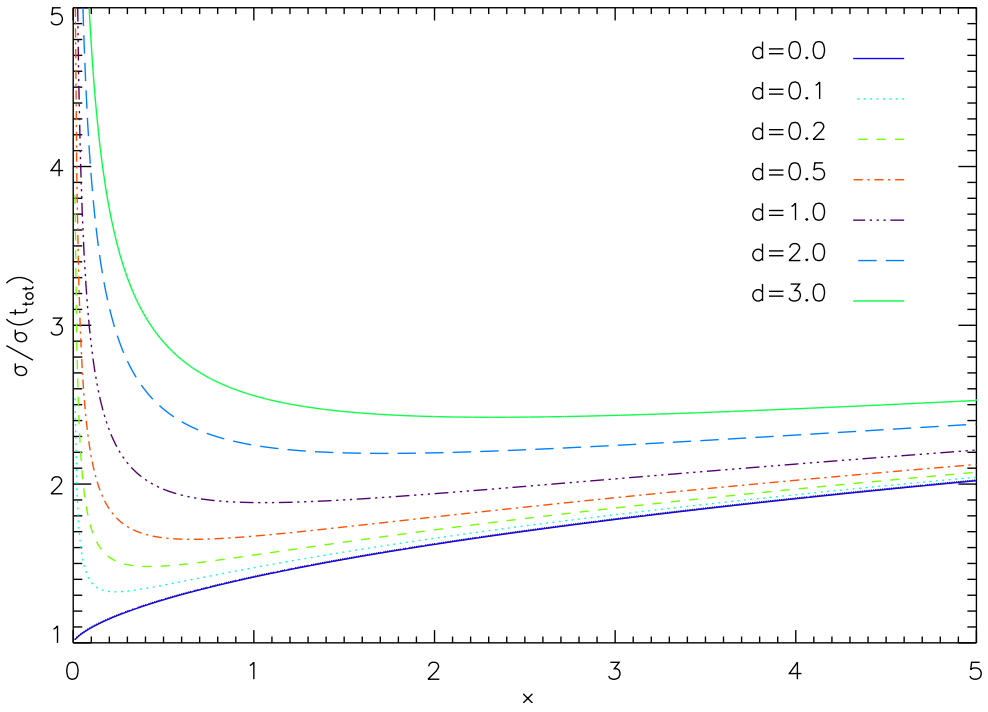

Figure 7: Total noise in chopped observations consisting of radiometric and drift contributions relative to the minimum radiometric noise as a function of the relative chop phase length x. The different curves represent different relative dead times per cycle. The drift noise is characterised here by a spectral index $\alpha = 0.7$.

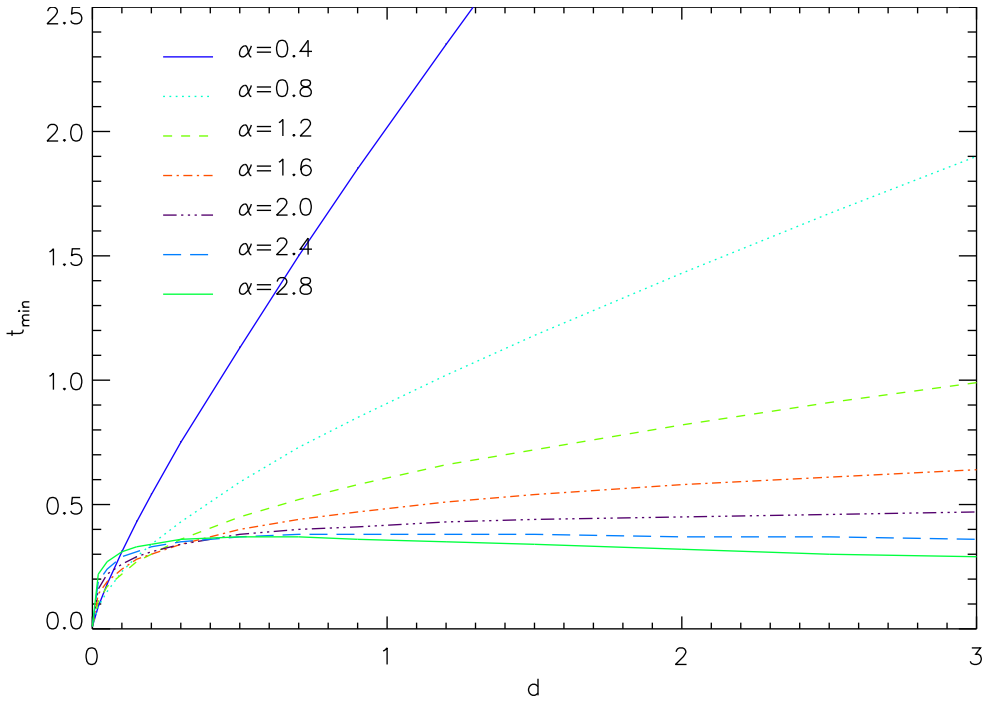

Figure 8: The optimum phase length depending on the relative overhead from the dead time for different spectral indices of the instrumental drift.

where root$_x\{\}$ denotes the root of the expression with respect to x.

Fig. 8 shows this optimum integration time providing the minimum total noise as a function of the dead time per cycle for a set of different spectral drift indices α. For spectral indices between 1.5 and 3 the curves hardly depend on the exact value of α. When the dead time exceeds half the Allan time the optimum integration time saturates also at about half the Allan time. For shallow fluctuation spectra, however, the optimum integration time increases strongly with the dead time so that it can easily exceed the Allan time for long dead times and spectral indices $\alpha < 1$.

This effect somewhat relaxes the constraints for planning observations aimed at an accurate measurement of the continuum level. Although the total-power Allan time is very short, the flat spectral index helps to lower the requirements with respect to the timing for switching between source and reference by allowing cycle times which exceed the Allan time.

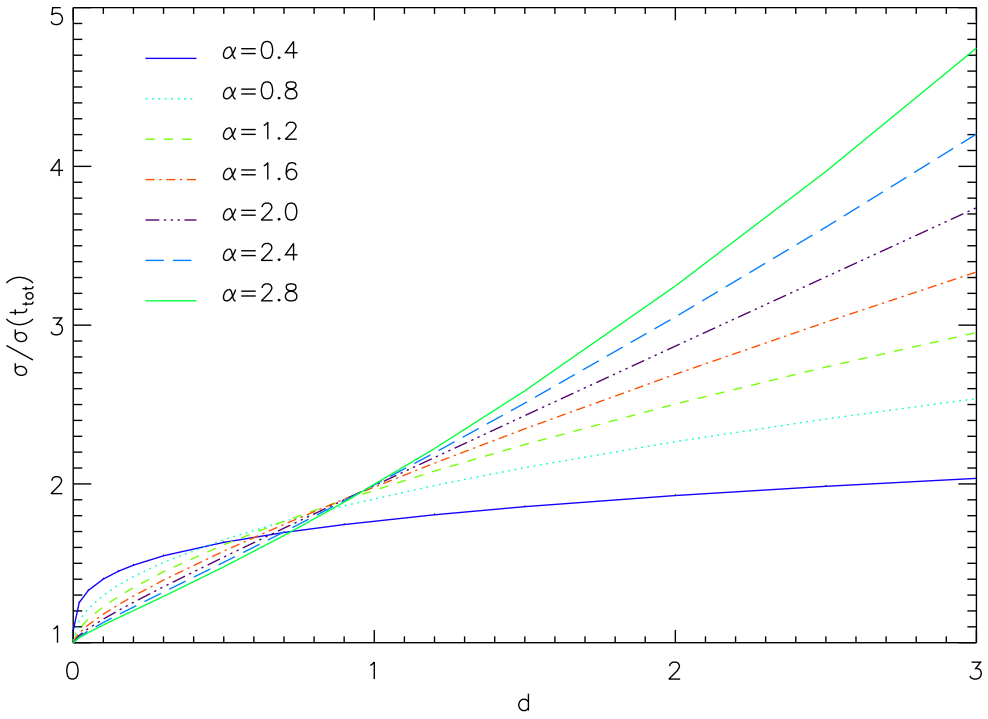

Figure 9: Total noise RMS in a chopped observations relative to the radiometric noise obtainable in the same observing time, $\sigma_{\mathrm{tot,opt}} \times \sqrt{B_{\mathrm{FI}} t_{\mathrm{tot}}}/(2\langle s \rangle)$, for the optimum phase length as a function of the relative dead time per cycle.

For the planning of the observations it is important to have a advance knowledge of the total uncertainty of the data which will be obtained in a certain observing time when knowing the properties of the instrument. Assumed that all observations can be performed in an optimum cycle time so that the total noise is minimised we can compute the total noise relative to the minimum radiometric noise like in Fig. 7 as a function of the spectral index of the instrumental drift and relative dead time. The result is shown in Fig. 9. For dead times exceeding about half the Allan time the drift contribution has the smallest impact for very shallow noise spectra whereas steep spectra naturally create drift contributions that grow quickly in time. As long as the dead time per cycle stays below the Allan time, the total noise RMS exceeds the noise RMS of an ideal instrument

only by a factor of up to two.

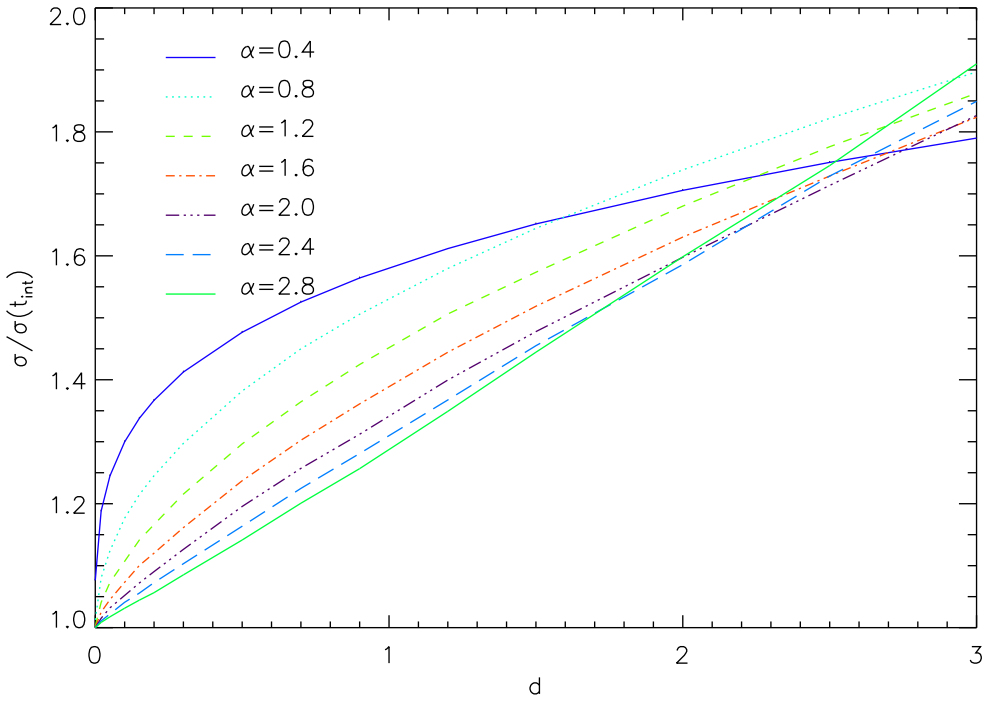

Figure 10: Drift noise relative to the radiometric noise for the optimum phase length as a function of the relative dead time per cycle.

To check which part of this increased noise is actually radiometric noise increased by the overhead from the dead time in each cycle and which parts stems from the instrumental drift one can also compare the drift noise relative to the radiometric noise of the actual observation, i.e. the quantity $\sigma_{\mathrm{tot,opt}}/\langle s\rangle \times \sqrt{B_{\mathrm{Fl}} t_{\mathrm{tot}}/(4+2d/x)}$. This is shown in Fig. 10. The ratio of the drift noise to the actual radiometric noise is a measure for size of possible ripples due to drift effects which might be visible in the noise of the baseline. Here, we see that in case of shallow drift spectra already for relatively small dead times a noticeably drift contribution is to be expected whereas steep spectra result in an almost linear growth of the drift contribution with dead time. This plot underlines that the total accuracy of an observation may not be estimated just by a fit of the RMS noise level in the data but that possible drift contributions always have to be taken into account.

6 Conclusions

6.1 The optimum method

We propose a new scheme for the computation of the Allan variance of a time series of spectrometer data. It combines the advantages of the spectroscopic Allan variance by Schieder et al. (1985) with the advantages of the baseline Allan variance by Siebertz (1998) and allows to analyse the stability of an instrument with respect to total-power fluctuations and with respect to differential fluctuations across the spectrum within the same mathematical scheme.

We propose two possible implementations for the algorithm to actually compute the Allan variance spectra which differ in the required computing power and the subjective "smoothness" of the resulting spectra although being both accurate within the achievable uncertainty of the total Allan variance analysis. In case of a similar behaviour of all channels across the spectrometer it is possible to characterise the instrument by a single Allan variance plot, but this has to be checked in each case by visualising the channel-by channel Allan variance.

A new definition of the instrument stability time allows to characterise the instrument also in case of $1/f$ or shallower fluctuation spectra but at the cost of not being directly comparable to the traditional Allan minimum time. Using this stability time we compute the optimum observing strategy for arbitrary spectral drift characteristics. It provides the minimum total uncertainty of the measured spectra per given observing time by adjusting the cycle time for switching between source and reference observations. We find that steep fluctuation spectra always help to reduce the total noise of an observation in case of small dead times needed for switching between source to reference, but that shallow spectra are advantageous in case of long dead times. Uncertainties due to drift noise always may show up as baseline ripples but the accuracy of an observation as a whole may only be judged when adding radiometric and drift noise.

6.2 The expected instrument stability for HIFI

The analysis of first test measurements with the high-electron-mobility transistors from the Centro Astronomico de Yebes and the Cologne wide-band spectrometer planned to be used for HIFI show that the total-power stability and the spectroscopic stability of the instrument will considerably differ. The total power Allan times are shorter by a factor 100 and the total power fluctuations show a spectrum shallower than $1/f$ noise. The gain fluctuations in the HEMTs used to amplify the signal from the mixers have a spectral coefficient of about 0.7. In contrast, spectroscopically normalised data where a zeroth order baseline is subtracted show fluctuations with a spectral index of about 2.

Although the formalism derived above allows to optimise both total-power and spectroscopic observations, the drastically lower stability of the instrument with respect to total power fluctuations shows that extremely short chop cycles are required to obtain a determination of the continuum level of a source without being limited by drift noise. If the instrument does not allow to use the optimum chop cycle the drift noise can exceed the radiometric noise by several orders of magnitude.

Acknowledgements

I want to thank Rudolf Schieder and Oliver Siebertz for direct help in the computations and many useful discussions. This work was supported by the DLR grant 50 OF 0006.

References

Allan D.W. 1966, Proc. IEEE 54, no. 2, 221

Kooi J., Chattopadhyay G., Thileman M., Phillips T., Schieder R. 2000, Int. J. Infrared & Millim. Waves 8, 689

Kraus J. 1980, Radio astronomy, Cygnus Quasar Books, Chapter 7
Roelfsema P. et al. 2002, "HIFI Calibration Use Cases", ICC/2001-005, 02/20/02
Schieder R., Rau G., Vowinkel B. 1985, in Instrumentation for Submillimetrer Spectroscopy, ed. E. Kollberg (Academic Press), 189
Schieder R., Kramer C. 2001, A&A 373, 746
Schieder R., Tolls V., Winnewisser G. 1998, Experimental Astronomy 1, no. 2, 101
Siebertz O. 1998, "Akusto-optisches Spektrometer mit variabler Auflsung", Ph.D. thesis, University Cologne
Whyborn N.D. 2003, DM IF-1 Amplifier Gain Stability Summary, SRON Technical Note 06/02/2003