"Sideband Smear": Single Sideband Discrimination in Double Sideband Receiver Systems

J.M. Payne, D.T. Emerson, P.R. Jewell and R.N. Martin

NRAO, 949 N. Cherry Ave,
Campus Bldg 65,
Tucson, Az 85718.

7 May, 1991

Abstract

In mm-wave astronomy the most common receiver is a cooled low noise double sideband mixer, feeding a low noise i.f. amplifier. Although it is possible to obtain some rejection of the unwanted sideband by adjusting mixer backshorts or by using a wave-optic sideband filter in front of the receiver feed, this is often not possible. Depending on the choice of intermediate frequency, adjustable backshorts may give only a limited degree of unwanted sideband rejection, and the receiver tuning for best sideband ratio may not give the optimum noise performance. A filter before the receiver feed results in some degradation of system noise, and additional complexity in adjusting the receiver for optimum performance.

The technique of sideband discrimination described here, known as "Sideband Smear," requires that the first local oscillator (lo) and the second or later lo be under computer control; this is normally the case in modern radio telescopes. Beyond this, no extra hardware is required. The system can discriminate between narrow-band emission occurring in the upper and lower sidebands of the first mixer, giving a rejection ratio of up to ~100, although noise input from the unwanted sideband, e.g. atmospheric or spillover radiation - is not attenuated.

Introduction

Nearly all receivers currently available for spectral line observations in the short mm-wavelength (~< 1mm) radio region are mixers. Most often, the receiver has nearly equal response to both upper and lower sidebands at the receiver feed. No preamplifiers have yet been made to work at these frequencies, and the extra loss and complexity of a filter able to reject the unwanted sideband at the input frequency is often undesirable. Backshort tuning elements in the mixer circuit itself can give significant attenuation of the unwanted sideband, but the use of mechanically adjustable tuning elements at cryogenic temperatures can give reliability problems. The recent development of very low noise, fixed tuned mixers that have equal response in both sidebands over a wide range of frequencies makes the sideband rejection method described in this note especially attractive.[1] A simple block diagram of a mm-wave radio telescope receiver system is illustrated in Figure 1.

Figure 2 shows an example spectrum taken at 23???? GHz, showing spectral lines occurring in both upper and lower sidebands of the input frequency. Fig 2(b) illustrates the same spectrum observed with a +10 MHz shift in local oscillator frequency. Features belonging to the different sidebands have moved in opposite directions, enabling an astronomer to interpret which features belong to which sideband. Of course this uses twice as much telescope observing time, and in a complex region there may still be some ambiguities as to which sideband certain features belong.

Most receiver systems involve more than one frequency conversion. Figure 3 shows a more typical receiver block diagram. Both local oscillators are under computer control.

Sideband Smear

The technique of shifting the first local oscillator to identify which features belong to which sideband is a special case of a more general technique. Figure 4 (a) illustrates a model spectrum, with an USB and a LSB feature. In Figure 4 (b) a first local oscillator shift of +10 MHz has been applied, shifting the USB feature to the left, and the LSB to the right. In Figure 4 (c) an offset of 10 MHz has been applied to the second local oscillator, shifting the LSB feature back to its original position (cf Fig 4(a)), but putting the USB feature even further from its original position in the spectrometer. If the observation simulated in Fig 4(c) is repeated, but with many (=N) different frequency shifts, the energy from the unwanted (USB) feature will be spread into the N different positions, each reduced in amplitude by (1/N). The wanted, LSB feature is unaffected, while the unwanted, USB feature has effectively been convolved with the frequency-step function, in this case consisting of N equal delta functions. By varying the number of coupled frequency steps of the first and second local oscillators, and adjusting the integration time spent at each frequency offset, the features in the unwanted sideband may be convolved with almost any chosen, positive, convolution function. An obvious example is obtained by using a large number of offset steps, each closely spaced with equal integration times; this becomes equivalent to a gradual, synchronized linear shift of frequency with time of the two local oscillators. This would be equivalent to convolving the unwanted sideband signals with a top hat function, of width equal to the total offset frequency excursion. Any feature from the unwanted sideband appearing in the resultant spectrum will have been smeared by the width of this function, and will appear as a slightly raised baseline to the features from the wanted sideband. Other functions are possible and may be advantageous in a given set of circumstances, such as a gaussian or a sinusoidal smear function.

By reversing the sign of the frequency offset of one local oscillators, alternate sidebands will be accepted and rejected. Note that random receiver fluctuations, which are not correlated in the different sideband smear offset steps, will not be smeared. Noise from the unwanted sideband will still degrade the signal-to-noise ratio of the wanted sideband in the same way as without Sideband Smear in operation.

This process is illustrated by the simulations shown in Figures 4 (a) to (e). A real example of the technique, using the NRAO 12m radio telescope at Kitt Peak, is shown in Figure 5.

Practical implementation of Sideband Smear

Although a very simple technique, there are some practical points that deserve attention:

1. Total range of offset frequency sweep. If the total i.f. bandwidth is not sufficient, there is the risk of putting part of the spectrometer beyond the i.f. response. If the width of spectrum of interest is B, the total bandwidth of the i.f. amplifier BI, then the total range of Sideband Smear frequency offset S should be:

$$S < (BI - B)$$

There may also be other constraints,- e.g. the first local oscillator is normally phaselocked in some way to a signal derived from a computer controlled frequency synthesizer of much lower frequency. The total range of frequency swing may be limited by the performance of the phase lock circuity and the high frequency oscillating device.

Subject to the above constraints, the larger the total frequency sweep, the bigger the discrimination, in terms of spectral width and amplitude, between the wanted and unwanted sideband.

2. Step size. If a series of small frequency offset steps is used, rather than a continuous sweep in frequency, then the offset frequency step should be small enough. For a spectrometer that fully samples the spectrum, such as a digital autocorrelator spectrometer, an offset step equal to the spectrometer frequency sampling interval, or less, is appropriate. { For a spectrometer that undersamples the spectrum, such as a filter bank with (channel spacing)=(channel width), then a step size either exactly equal to the channel spacing, or a two or three times more closely spaced, is required. [delete last sentence?]}

3. Step timing. In general, the faster the offset frequency is stepped, the better. The total frequency sweep (see (1) above) has to be covered within the total integration time of a given observation on the sky. However, there may be limitations due to computer overhead and frequency synthesizer or phase lock settling time. A reasonable compromise might be a few steps per second.

4. Signal and Reference matching. Most observing techniques involve switching in some sequence between the wanted "signal" position on the sky, and a blank region known as the "reference" position. It is most important that the sequence of Sideband Smear offset frequency
steps match in the signal and reference observations. That is, the pre-programmed frequency offset step sequence should restart at the same point for signal and reference observations. If this is not done, then bad spectral baselines, due to the inevitable frequency structure in the telescope optics, receiver frontend and if passband, are likely to appear in the spectrum.

A possible alternative to this carefully synchronized switching sequence would be a very rapid, asynchronous frequency sweep control. The first and second local oscillators would have swept the entire smear range many times - ideally many hundred times - synchronously with each other, but asynchronous to the start of data acquisition, within the total integration time.

Possible extension

Figure 6 shows an efficient extension to this technique, in which both sidebands are retrieved independently, but at the cost of duplicating the second mixer and spectrometer. Upper and lower sideband signals are retrieved independently, with the opposite sideband being smeared in each spectrum.

Conclusion

Many short mm-wave radio telescopes use double sideband mixers, and cannot easily distinguish between upper and lower sideband signals. We have described a technique for separating lower and upper sideband signals in such systems. This technique involves no extra hardware, involving relatively easy software programming of existing local oscillator frequency synthesizers. We have illustrated the technique with results from the NRAO 12m radio telescope at Kitt Peak, where the technique is now in regular use {!}.

References

1.A.R. Kerr, S-K Pan, "Some Recent Developments in the Design of SIS Mixers," *International Journal of Infrared and Millimeter Waves*, Vol. 11, No. 10, 1990.

Figure Captions

Figure 1:

A simplified block diagram of a mm-wave spectral line receiver system. The incoming signal is combined with a local oscillator in a low noise mixer device, then further amplified in the intermediate amplifier strip. Example frequencies are shown; with a 230 GHz local oscillator and a 1.5 GHz i.f., both representative values, the system is sensitive to 500 MHz bands centred both on 231.5 and 228.5 GHz, the upper and lower sidebands. The output of the i.f. amplifier is fed to a spectrometer. The spectrometer might have, for example, 250 separate filters, each of 2 MHz bandwidth and separated by 2 MHz, covering the entire 500 MHz i.f. bandwidth. Each filter is followed by a square law detector and integrator, recording the average power in some specified integration time in each part of the spectrum.

Figure 2:

(a) An example spectrum showing a 500 MHz (?) band centred on 230.5 (?) GHz, in the direction of the Orion Nebula. The many spectral features come from a variety of molecular transitions. Some of these are received in the upper sideband (232 - 232.5 GHz (?)) and some from the lower sideband (229 - 228.5 GHz {?}).

(b) A standard technique used for astronomers to separate out USB emission from LSB signals is to make a second observations, but shifting the local oscillator by a small amount - e.g. +10 MHz. The USB spectral lines will then be shifted by -10 MHz in the i.f. chain, while the LSB lines will be shifted by +10 MHz relative to the first spectrum. Inspection of the two spectra enables the astronomer to decide which features are USB and which LSB, albeit with some small ambiguities. Figure (b) shows the same spectrum as (a), but with the first lo shifted from 230.?00 to 230.?10 GHz.

{Can we borrow a spectrum from Phil, Lucy or Barry, showing lots of lines in both sidebands?)

Figure 3:

A block diagram of a typical receiver setup used at mm-waves. There is no filter before the first mixer, which has a double sideband response. The spectrometer is a device for measuring the average power at a number of points within the i.f. passband. Typically, the total i.f. bandwidth might be 500 MHz, with the spectrometer measuring 250 independent 2 MHz bands within this passband.

Figure 4:

(a) A simulated observation showing two simple spectral lines, one in each sideband. The total width of the spectrum is 256 MHz, with one spectral channel per MHz.

(b) The simulated observation when the local oscillator frequency is changed by +10 MHz. The signal from the LSB is shifted to the right, while the USB line is shifted to the left by 10 MHz.

(c) The same simulated observation as (b), but with the second local oscillator shifted by +10 MHz. (The second mixer is assumed to be a single sideband mixer, with local oscillator lower in frequency than the first i.f.) It is seen that the original LSB line is now at the same spectral channel as the original observation (a), but that the USB signal has shifted 20 MHz in the spectrometer from the original observation (a).

(d) The average of 10 observations such as (c), but with coupled first and second local oscillator shifts of -20,-10,0,10, and 20 MHz. The LSB signal is unaffected, but the power from the USB signal has now been smeared over 5 separate positions.

(e) As (d) but using 256 different, coupled lo shifts, in steps of 1 MHz. The wanted, lower sideband signal is unaffected, but the unwanted, USB signal has been smeared over 256 points. An unwanted signal in the centre of the passband will be smeared over the entire passband, appearing as a small D.C. offset in the baseline of the spectrum. Although the total flux, brightness temperature integrated over frequency, of the unwanted line is the same, it has been reduced in peak amplitude by a factor of ~100 (=256 MHz/line width) and is insignificant. For a line initially at the edge of the passband, there will be a small discontinuity of (Peak temperature*(line width)/256 MHz) near the centre of the passband, but this too is insignificant. Any features narrower than 256 MHz in this spectrum must result from the desired, LSB signal. Unwanted signals in the upper sideband have been smeared over 256 MHz.

Note that random noise is uncorrelated between each 256 MHz lo step, and so will not be smeared in the way that true signals will. i.e. random noise from both sidebands will appear in the spectrum.

Figure 5:

A real example of the simulation illustrated in Figure 4. Strong emission from the Orion Nebula is observed, giving the strong line in (a) from the lower sideband. "Sideband Smear" is not in operation for this observation, with first and second lo fixed in frequency.

In (b) the telescope continues to track Orion, but "Sideband Smear" is in operation in the sense of rejecting lower sideband, leaving upper sideband signals intact.

In (c) the sense of "Sideband Smear" is such that the first and lo track correctly for lower sideband, but will smear out any upper sideband signal. The strong line from Orion is now identical to the original observation of (a).

Figure 6.

A possible extension to the Sideband Smear technique, aimed at making more efficient use of telescope observing time. The first local oscillator steps or sweeps frequency in the normal way, but the second local oscillator, and spectrometer, are duplicated. In one case the second oscillator (a) sweeps in synchronization with the first oscillator increasing in frequency, but second oscillator (b) sweeps in decreasing frequency. In one spectrometer (a) the USB signal is retrieved with the LSB signal smeared, while in (b) the spectrometer integrates the LSB signal without distortion, while the USB signal is smeared. This doubles the data acquisition efficiency of the telescope.