EVLA Memo 110
The Effect of Amplifier Compression by Narrowband RFI on Radio Interferometer Imaging

Rick Perley and Bob Hayward

April 25, 2007

Abstract

An experiment is described which has permitted direct measurement of the effects from coupling out-of-band emission through the third-order intermodulation product into the astronomical band of interest, as a function of the degree of compression of a front-end amplifier. The coupling was measured by injecting a 1545 MHz tone into four of the VLA's FE amplifiers, and measuring the phase and amplitude of the 1665 MHz OH maser at its 'reflected' frequency of 1425 MHz, as the saturating tone power was varied.

We find the following effects:

- The effective gain of the compressed amplifiers for astronomical signals is reduced by about twice the compression level.
- The correlation coefficient of in-band emission is reduced by ∼5% at 1 dB compression, and by ∼20% at 3 dB compression, due to the addition of incoherent out-of-band emission.
- Out-of-band emission is reflected into the band of interest at a level of about 1% of its natural strength at 1 dB compression, and ∼25% at 3 dB compression.
- The phase of the reflected emission rotates at typically twice the natural fringe rate, relative to the fringe-stopped emission in the observing band.

The rapid phase rotation of the reflected emission will in general provide considerable attenuation of these signals in interferometric imaging.

1 Introduction

Radio interferometers currently in planning or under construction, particularly those working in the meter and centimeter wavelength bands, are designed with very wide bandwidths – often exceeding 2:1 in bandwidth ratio – to maximize continuum sensitivity and to enable unrestricted frequency access for spectral line observations. These wide bandwidths will however permit strong radio-frequency interference (RFI) into the amplification and signal transmission systems, with the likelihood of amplifier compression and non-linear response. Although good engineering design can provide very high linearity, there will always be some signals entering the system which will cause, at least momentarily, compression in the analog signal system. The consequences of observing with amplifiers operating in their non-linear regimes can be serious – power in the fundamental is shifted to higher harmonics, causing an effective loss of gain, and a host of intermodulation products which can shift out-of-band information into the astronomical band of interest. The latter of these is the most serious for interferometric imaging, and is the subject of this memorandum.

Our first attempt to study the consequences of compression [1] utilized the high temperature broadband solar calibration noise diodes contained within each antenna. These turned out to be not strong enough to set the system far enough into compression to cause any notable effects. In any event, the use of broadband signals to saturate the receivers is not a good model for the anticipated operating environment, where strong quasi-CW signals are responsible for the compression.

Our second attempt [2] utilized a strong CW tone to drive the C-band receivers of four VLA antennas to compression levels between 1 and 6 dB. This experiment was successful in determining the modest loss of SNR due to compression, and showed that even with these levels of compression, there was no measureable loss of closure – in essence, the astronomical information contained within the coherence function was preserved, albeit with lower sensitivity.

Unfortunately, due to the setup chosen for that experiment, the 3$^{\text{rd}}$-order intermodulation product itself was filtered out, and we were unable to measure the amplitude and phase of the reflected signal in the observing band. This filtering, although desirable in the sense of keeping unwanted reflected information out of the band of interest, will not occur in practice, as the entire front-end bandpass will normally be present at the saturating element in the analog receiver chain.

It is important both for the design of wide-bandwidth radio telescope receivers, and for deciding on mitigating strategies, that the strength of these third-order products, and their ramification on imaging be understood. In this, our third (and we hope last!) study, we report the results of an experiment specifically designed to measure the 3$^{\text{rd}}$order harmonic coupling due to amplifier gain compression.

2 Non-Linear Characteristics of Amplifiers in Compression

In [2] we presented a short mathematical description of how amplifier compression will shift information from one frequency to another by non-linear coupling. We repeat and extend that analysis here.

The essence of the problem inherent in non-linear devices is shown in Figure 1. The left panel shows the normalised instantaneous voltage transfer functions for a perfect amplifier (black line), and for a model of a real device, (in red), for which the output voltage asymptotes to a maximum value, $V_0 = G_v V_s$ (dashed blue) as the input voltage, V_i exceeds its 'saturation level', V_s (dashed green). We define V_s as the input voltage at which, for an ideal linear amplifier, the output voltage equals the limiting voltage. To illustrate the effect of saturation on the output voltage, the right panel shows the output voltage profiles for three different inputs: (a) a low amplitude signal (black), for which $V_i = 0.25 V_s \cos(\omega t)$; (b) a slightly saturated signal (red), $V_i = V_s \cos(\omega t)$, for which the maximum output amplitude is $\sim 0.75 G_v V_s$; and (c) a high amplitude signal (blue), $V_i = 5V_s \cos(\omega t)$, for which the ouput is strongly limited at $\sim G_v V_s$. The effect of a limiting output voltage is to 'square off' the peaks in the output amplitude, resulting in conversion of power from the fundamental frequency into higher harmonics.

A better understanding of the consequences of harmonic conversion can be gained with a straightforward mathematical analysis utilizing an idealized transfer function given by

$$V_o = G_v V_s \tanh\left(\frac{V_i}{V_s}\right), \tag{1}$$

where G_v is the voltage gain, and V_s is the input saturation voltage. This arctangent representation seems to provide an excellent description of saturation [3] and has convenient mathematical properties.

We consider the harmonic content of the output signal when a pure harmonic signal, $V_i(t) = A\cos(\omega t)$, is input. With our amplifier model, the output signal is given by

$$V_o = G_v V_s \tanh\left[\frac{A\cos(\omega t)}{V_s}\right] = \frac{G_v A}{X}\tanh[X\cos(\omega t)], \tag{2}$$

where $X = A/V_s$ is the normalised input amplitude. We are interested in the output response when the amplifier is slightly saturated – when $X \leq 1$. In this case, we can expand the hyperbolic tangent

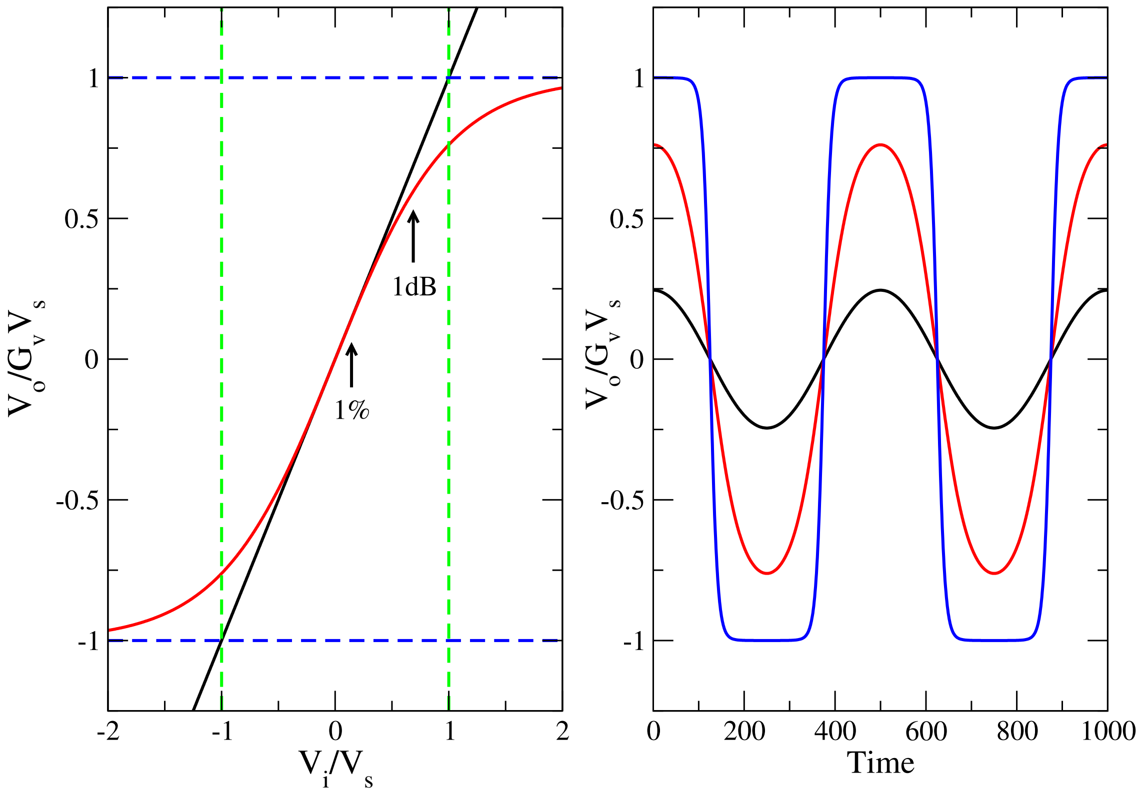

Figure 1: Left Panel: The black line shows the instantaneous voltage transfer function of an ideal, linear, amplifier. The red line shows an idealization of a real amplifier, for which there exists an asymptotic limiting voltage V_o shown by the dashed blue lines. The green dashed lines define the saturation voltage, V_s, defined by $V_o = G_v V_s$. Also shown are the input amplitudes corresponding to the 1% and 1dB compression levels. Right Panel: Oscillographs of three voltage outputs, demonstrating the responses to a non-saturating sinusoidal input (black) $V_o = 0.25 V_s \cos(\omega t)$, slightly saturating input (red), $V_i = V_s \cos(\omega t)$, and heavily saturating input (blue), $V_i = 5 V_s \cos(\omega t)$.

to find, for the normalized output voltage,

$$\frac{V_0}{G_v A} = \cos(\omega t) - \frac{X^2}{3}\cos^3(\omega t) + \frac{2X^4}{15}\cos^5(\omega t) - \frac{17X^6}{315}\cos^7(\omega t) + \cdots \tag{3}$$

The series formally converges for $|X| < \pi/2$.

We are interested in the power in the various harmonics of the fundamental. These can be determined from expanding the odd powers of the cosine function into its harmonic components[1] which allows us to write the normalized output amplitude as

$$\frac{V_o}{G_v A} = F_1 \cos(\omega t) + F_3 \cos(3\omega t) + F_5 \cos(5\omega t) + \cdots \tag{4}$$

where the amplitude coefficients, F_n, are

$$F_1 = 1 - \frac{X^2}{4} + \frac{X^4}{12} - \frac{17X^6}{576} + \cdots \tag{5}$$

[1]conveniently done, using the relation, valid for n odd,

$$\cos^n(\theta) = \frac{1}{2^{n-1}} \sum_{k=0}^{(n-1)/2} \binom{n}{k} \cos(n-2k)\theta$$

$$F_3 = -\frac{X^2}{12} + \frac{X^4}{24} - \frac{119X^6}{1696} + \cdots \quad (6)$$

$$F_5 = \frac{X^4}{120} - \frac{17X^6}{360} + \cdots \quad (7)$$

The power within each component is found from squaring these coefficients. This exercise shows that the spectral content of the output voltage is composed entirely of odd harmonics of the fundamental input frequency.

The compression C is defined by the ratio of the power in the fundamental harmonic of the output signal $(G_v A F_1)^2$, divided by that provided by an ideal, linear amplifier of the same low-voltage gain, $(G_v A)^2$. Thus, we find

$$C = \left[1 - \frac{X^2}{4} + \frac{X^4}{12} - \frac{17X^6}{576} + \cdots\right]^2 \quad (8)$$

This equation can be solved algebraicly to find that at 1% compression ($C = 0.99$), $A_{1\%} = 0.142V_s$, and at 1 dB compression, ($C = 0.7943$), $A_{1dB} = 0.708V_s$. The input power level at the 1% compression level is 13.9 dB below that of the 1 dB compression level.[2].

The expressions derived above can only be used for very small compression, due to the very slow convergence of the hyperbolic tangent polynomial. To estimate the power in the harmonics for arbitrary degrees of saturation, we have utilized a DFT to harmonically decompose a sample of the output voltage, with the results shown in Figure 2. This shows the normalized power output for an ideal amplifier (black), and from our model amplifier (red), as a function of the input power. As the amplifier goes into compression, the output power in the fundamental tone falls below the idealized extrapolation of the low-power response. The green and blue solid lines show the power in the third and fifth order harmonics. These also curve to a constant power as the amplifier goes into higher compression, and more and more power is converted to higher harmonics. The low-power asymptotes of the third and fifth order harmonics are shown in dashed green and blue, and are described by the expressions given in equations 6 and 7, respectively. The interception points between these extrapolations and the idealized output power in black are referred to as the $3^{\rm rd}$and $5^{\rm th}$harmonic intercepts.

3 Intermodulations from Two Input Tones

A single tone strong enough to cause compression is not in itself a serious concern. The output consists of odd-numbered harmonics of the fundamental, and in nearly all cases the overtones will be filtered out before reaching the sampler, leaving an attenuated, but uncontaminated, version of the fundamental. The problems arise when the strong interfering signal interacts with other signals within the band of interest. To investigate the effect of two signals entering a non-linear device, we input to our model a sinusoidal input given by

$$V_i = A\cos(\omega_a t + \phi_a) + B\cos(\omega_b t + \phi_b) \quad (9)$$

The output is then given by

$$V_o = G_v V_s \tanh\left[X\cos(x) + Y\cos(y)\right], \quad (10)$$

where $X = A/V_s$ and $Y = B/V_s$ are the normalized amplitudes of the input signals, and $x = \omega_a t + \phi_a$ and $y = \omega_b t + \phi_b$ are the instantaneous phases. Expanding the hyperbolic tangent as before, and

[2]These relations permit alternate ways of expressing the voltage transfer function. For example, in terms of the 1 dB compression amplitude V_1, we can write $V_0 = 1.412 G_v V_1 \tanh[0.708 A\cos(\omega t)/V_1]$.

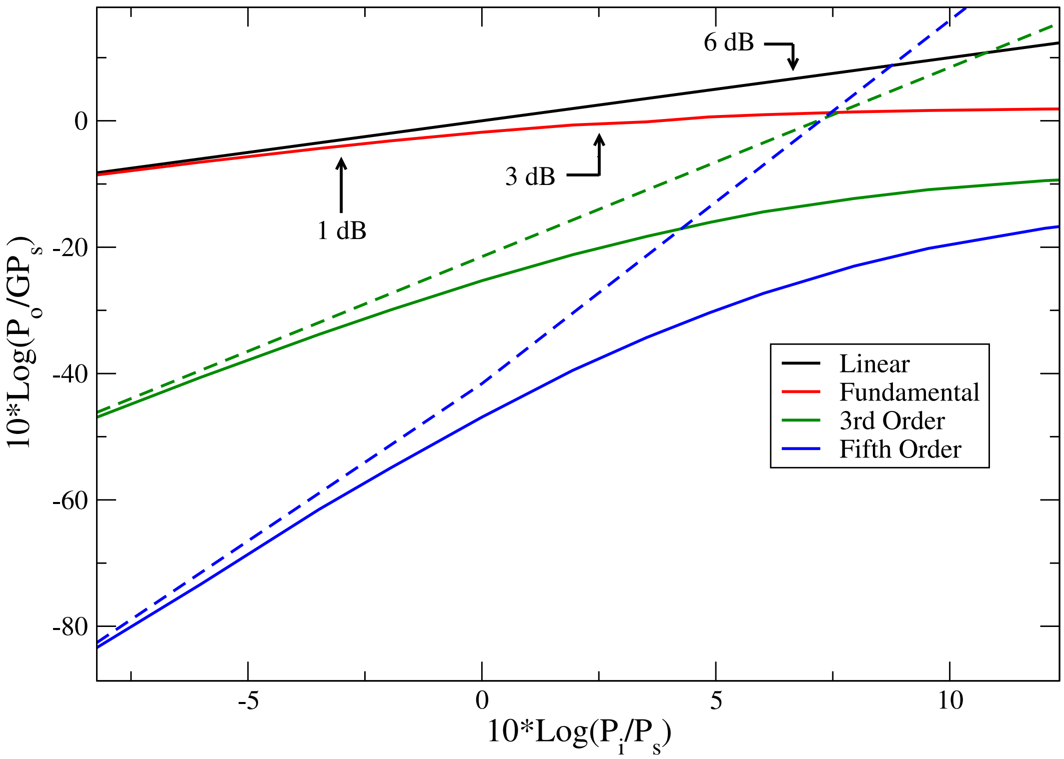

Figure 2: The power in the fundamental, 3^{rd}-order, and 5^{th}-order harmonics for our idealized amplifier response. Marked are the 1, 3, and 6 dB compression points, where the output power in the fundamental harmonic is reduced by a factor of 0.795, .501, and .251, respectively from that of a purely linear device. The 1% compression point is located -13.9 dB below the 1 dB point. The intercepts between the low-power extrapolations shown in dashes and the linear response in black are the $3^{rd}$and 5^{th}order intercepts.

harmonically decomposing to the third order only, we find

$$\begin{aligned}\frac{V_0}{G_v} &= A\left(1-\frac{X^2}{4}-\frac{Y^2}{2}\right)\cos x + B\left(1-\frac{Y^2}{4}-\frac{X^2}{2}\right)\cos y - A\left(\frac{X^2}{12}\right)\cos(3x) - B\left(\frac{Y^2}{12}\right)\cos(3y) \\ &\quad -A\left(\frac{XY}{4}\right)[\cos(2x-y)+\cos(2x+y)] - B\left(\frac{XY}{4}\right)[\cos(2y-x)+\cos(2y+x)] \qquad (11)\end{aligned}$$

In addition to the two fundamental and two third-order harmonics, we see the presence of four intermodulation products – terms at frequencies given by $\omega = 2\omega_a - \omega_b$, $2\omega_b - \omega_a$, $2\omega_a + \omega_b$, and $2\omega_b + \omega_a$. The two difference terms are of special concern, as they will normally lie close to the two input signal frequencies, and hence can lie within the system passband. The other two (summed) intermodulation products can normally be ignored as they will lie outside the bandwidth of the system response.

Extending the analysis to include 5^{th}order products adds the two fith-order harmonics, and eight new intermodulation products to the output spectrum. As an amplifier goes deeper into compression, every pair of input signals will produce tens or even hundreds of intermodulation products, many of which will appear within the desired bandpass.

In general, we note that the strength of the third-order intermodulation products rises as A^2B, or B^2A, whereas the power in the fundamentals rises slower than with A or B. Thus, the ratio of the power in the intermodulation products to that of the fundamental signals rapidly increases with increasing amplifier compression.

We now consider the case where the 'Y' input signal, representing the RFI, is strong enough to cause gain compression, and the 'X' signal, representing the astronomical information, is very weak. Thus, we assume $Y \gg X$, (so that $B \gg A$), that $X \ll 1$, and that $Y \leq 1$. Dropping terms much less than one, we can then write the output signal, normalized to the idealized output amplitude of A, as

$$\frac{V_0}{G_v A} = \left(1 - \frac{Y^2}{2}\right)\cos x + R\left(1 - \frac{Y^2}{4}\right)\cos y - R\left(\frac{Y^2}{12}\right)\cos(3y) - R\left(\frac{XY}{4}\right)[\cos(2y - x) + \cos(2y + x)] \tag{12}$$

where $R = B/A$ is the amplitude ratio between the interference signal and the astronomical signal. This formulation allows us to compare the strength of the output harmonics to that of the desired (fundamental) output. Because of the approximations involved, the terms shown are the low-power approximations of the actual relations. These expressions should not be employed for compression levels exceeding 1 dB.

To obtain the amplitudes of all output spectral components over a wide range of input RFI powers, we have employed a DFT routine to harmonically decompose the signal described in Eqn. 10. The results are shown in Figure 3. In this simulation, there were two tones submitted to the model amplifier.

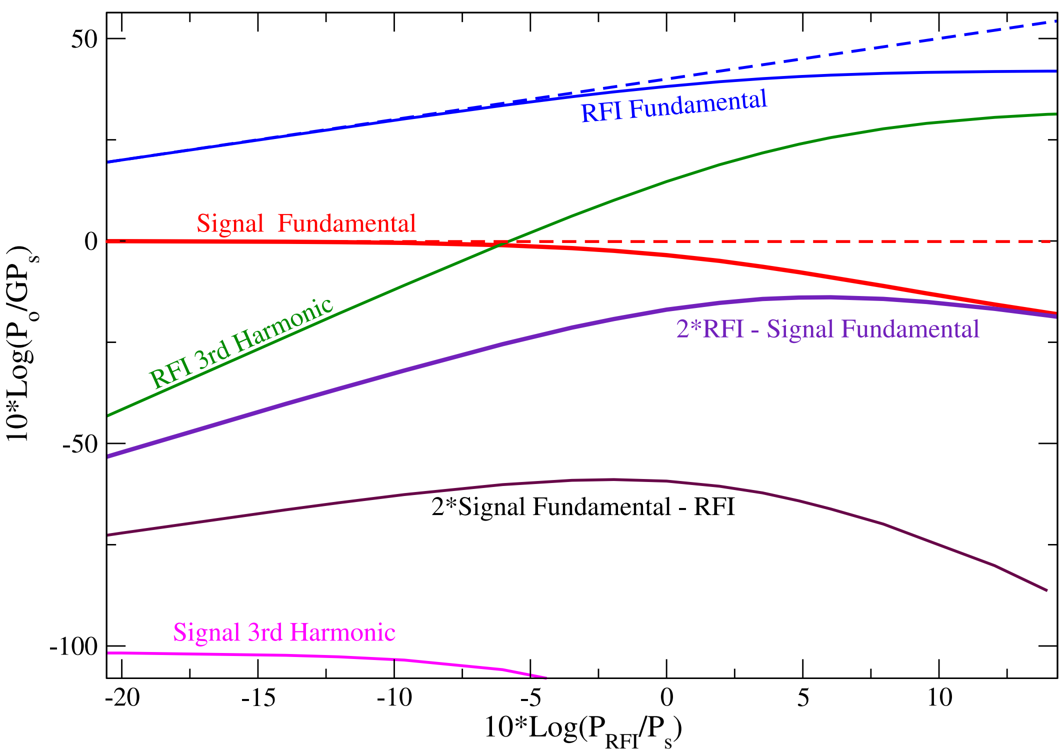

Figure 3: Showing all the fundamental, 3rd-order harmonics, and 3rd-order intermodulation products for our amplifier model, as a function of the power in the saturating RFI signal. The astronomical signal strength is fixed 40 dB below the saturation power. The blue line shows that the power in the fundamental of the saturating tone asymptotes to a constant value as the amplifier goes into compression. The red line shows the power in the (constant) astronomical signal, which declines as its power is shifted into higher harmonics. The purple line shows the rise of the third-order intermodulation product, whose power approaches that of the astronomical signal fundamental as saturation increases.

The 'signal tone' was fixed at a power level 40 dB below the saturation power, to ensure that it is not responsible for the amplifier compression, while the 'RFI tone' was varied from -40 dB below to +60 dB above the saturation power. The power in the various harmonics and intermodulation producrts are shown plotted as a function of the input RFI tone power. As the RFI tone power increases, and the amplifier goes into compression, the following occur:

- The output power in the RFI signal (blue) rises close to linearly, then asymptotes to a constant level as the amplifier goes into compression. The ratio between the extrapolation and output power defines the level of compression.

- The output power of the (constant input) astronomical signal (red) slowly declines, as the compression limits the range of its voltage, and converts fundamental harmonic power into overtones and intermodulation products. Notable is that the loss in signal power is twice that of the RFI power for low levels of compression.

- The power in the 3$^{\rm rd}$harmonic of the RFI (green) rises quickly, following the same relation shown in Fig. 2. As this tone will normally be filtered out before reaching subsequent stages of amplification or sampling, it is of no concern.

- The 3$^{\rm rd}$harmonic of the signal power (magenta) is very weak, and of no concern.

- The intermodulation products formed from twice the signal frequency and the RFI frequency ($2\omega_{sig} \pm \omega_{RFI}$) (brown) are also very weak, and of no consequence.

- The two intermodulation products formed from twice the RFI frequency and the signal frequency (purple) ($\omega_- = 2\omega_{RFI} - \omega_{sig}$, and $\omega_+ = 2\omega_{RFI} + \omega_{sig}$) rise quadratically, and approach the power of the astronomical signal itself as the amplifier goes into compression. The summed response will normally lie outside the passband, so is of no concern. The difference frequency will normally lie in the desired bandpass, and is the signal of concern.

For modest levels of compression, the fractional loss in signal power scales as $1-(1-Y^2/2)^2 \sim Y^2$, while the fractional loss in RFI (tone) power scales as $1-(1-Y^2/4)^2 \sim Y^2/2$. thus, the fractional loss in astronomical noise is twice that in tone power. Similarly, the signal attenuation, expressed in decibels, is twice that of the tone attenuation [3].

Hence, when the amplifier is at a 1% compression due to the RFI, the signal output is reduced by 2%, and when the compression is 1dB, the astronomical signal is reduced by $\sim$ 2dB. This behavior can be understood by noting that the weak astronomy signal voltage is added to that of the saturating RFI, and because of the non-linear nature of the transfer function, suffers a greater compression than that of the RFI. In the limit, when the RFI amplitude is so great that the output is a pure switching square wave, the power in the astronomy signal will be entirely lost, while that of the RFI is limited to that offered by the fundamental harmonic component of the output saturation level.

The relation between the astronomical signal power and the key intermodulation product is shown in detail in Figure 4.

Three important conclusions arise from this:

- Amplifier compression causes a significant reduction in astronomical signal power. The loss is by a negligible 2% at 1% compression, but reaches a significant 2 dB (27%) when the amplifier has reached a 1 dB compression. In some circumstances, this loss can be corrected, as discussed below.

- The coupling of the intermodulation product signal into the astronomical band is negligible at 1% compression levels, but becomes significant (2%) at 1 dB compression. Thus, at this level of compression, a strong out-of-band spectral line can appear in the band of interest, diminished by a factor of $\sim$ 50.

[3]If this is not immediately obvious, recall that $\ln(1+x) = x - x^2/2 + x^3/3 - \cdots$

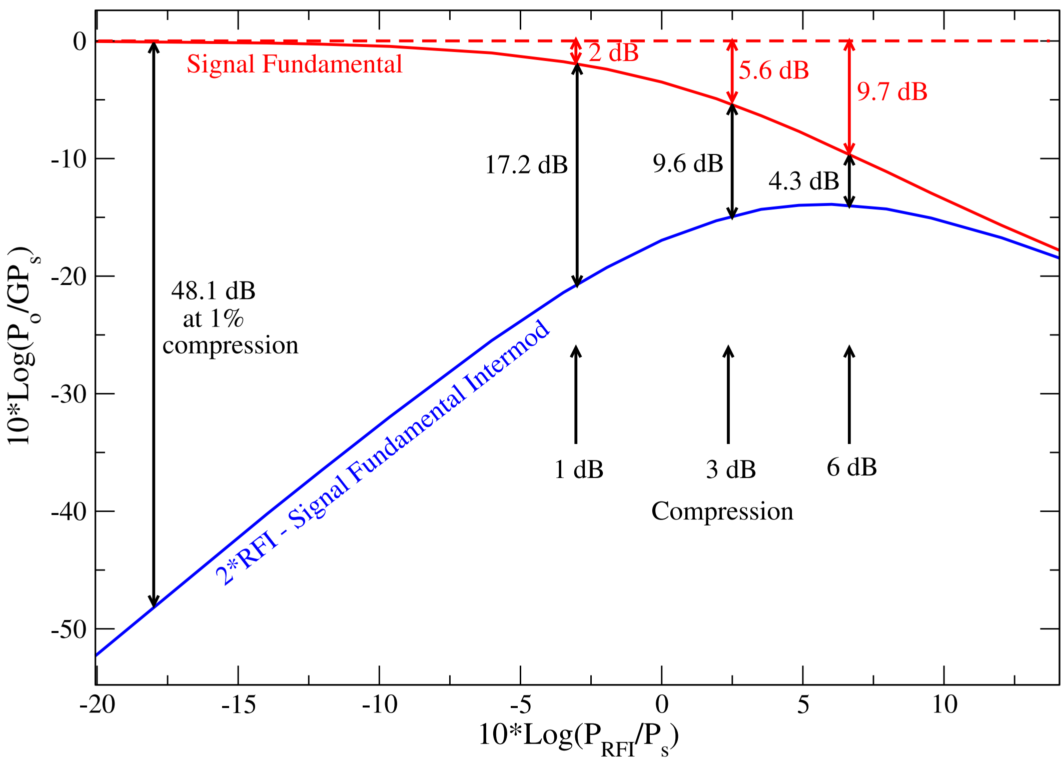

Figure 4: The power in the key intermodulation product (blue) approaches that in the astronomical signal as the RFI power rises to cause the amplifier to go into compression. With the amplifier at 1% compression, the power in the fundamental tone is reduced by ∼2%, while the intermodulation is down by a safe -48.1 dB. At 1 dB compression, the signal power is reduced by 2 dB, and the intermodulation product is now down by only -17.2 dB. The ratio between fundamental and intermodulation powers rapidly declines with increasing compression.

- With 3 dB compression, the out-of-band emission contribution reaches 50% of the in-band emission.

The loss in astronomical signal as shown by the red curve is effectively the same as a gain change in the amplifier. The astronomical signal, and any calibration signals introduced before the compressing amplifier are reduced nearly equally (provided the compression is not very high), and hence the loss in effective gain can be detected and corrected for provided the saturating signal is constant over the calibration time. There will be a loss in sensitivity involved, as the relative noise contribution of downstream electronics will increase. However, this loss in SNR will be small if the saturating component is located after the first stage of amplification. If the compressing signal is highly time variable, gain loss compensation will be set by the average loss over the integration time.

As the amplifier goes into compression due to an RFI tone, out-of-band emission from the 'image' frequency is added to that for the band of interest. Total power systems cannot discriminate between these two origins. In particular, if the switched noise power is present in the aliased band (as it normally will be), the synchronous demodulation of this calibration will not see any difference, with the result that the measured system temperature of the summed signal will not change due to the compression. For very high levels of compression, a second-order effect should be noted – the detected increment in power upon addition of the calibration noise will be less than for an unsaturated system, leading to a small apparent increase in the system temperature.

The next section considers applications to interferometry.

4 Interferometer Response

The preceding analysis considered only the amplitude and power responses due to amplifier compression. In interferometry, the relative phases of the signals are important, and we consider these effects in this section.

We consider the outputs of two antennas, both of which receive an astronomical signal of frequency ω_o. We take one of these antennas as the phase reference. Constant electronic phase shifts are ignored here, as it is the time dependency of the differential phase that we are interested in. Antenna 1's output is proportional to $\cos(\omega_o t)$. The second antenna's output is proportional to $\cos[\omega_o(t-\tau_g)]$, where τ_g is the geometric delay. Each signal then passes through an amplifier which has been put into compression by a signal of frequency ω_r. The third-order intermodulation signal for the first antenna has a time dependance of $\cos[(2\omega_r-\omega_o)t]$, while for the second antenna, it is given by $\cos[(2\omega_r-\omega_o)t-\omega_o\tau_g+2\phi_r]$, where ϕ_r is the phase of the RFI saturating signal at antenna 2, relative to that of antenna 1.

In order to maintain coherence, the signal of antenna 1 must be delayed by a time equal to the geometric delay applicable to antenna 2. Hence, the time dependency of the signal on antenna one following the insertion of the delay is: $\cos[(2\omega_r-\omega_o)(t-\tau_g)]$. The output from the correlator is given by the low frequency component of the product of the signals – the phase difference – from antennas one and two: $\cos(2\omega_r\tau_g+2\phi_r)$. The phase rate of the product is the observable of interest, and is given by

$$\dot{\phi} = 2\omega_r\dot{\tau}_g + 2\dot{\phi}_r \tag{13}$$

where $\dot{\tau}_g = -\omega_e B_u \cos\delta/c$ is the rate of change of delay, $\omega_e = 7.27\times10^{-5}$ rad/sec is the angular rotation rate of the earth, and B_u/c is the light-travel time of the E-W component of the baseline. This result is independent of whether the interferometer is 'direct', or employs a local oscillator conversion, with appropriate phase rotation.

The second term in the RHS of Eqn. 13 accounts for a phase differential in the saturating signal between stations. For stationary RFI, this term will be close to zero (non-zero being possible due to change in refraction, or antenna motion). For the experiment to be described in the next section, the saturating signal had a phase difference due to its origin.

We now consider the effect of the aliased signal on the correlation coefficient. The correlator does not measure the visibility amplitude directly, but rather estimates the correlation coefficient. This is converted to a visibility flux density by a correction factor dependent upon the system temperatures of the two antennas involved. We have argued above that the broadband noise aliased into the band of interest due to a strong RFI tone retains the switched noise used for estimation of the system temperature, so that no significant change in system temperature will be measured due to the amplifier compression. This is not the case for the correlation coefficient, as the aliased noise does not retain either the correct phase, nor the phase rate, of the emission in the band of interest. Hence, the measured correlation coefficient will be reduced by a factor dependent upon the coupling of the out-of-band emission into the band of interest. The measured correlation coefficient will be reduced as

$$\rho = \frac{\rho_0}{1+\epsilon} \tag{14}$$

where ρ_0 is the correlation coefficient for the band of interest, and ϵ is the fractional contribution of the aliased signal to the fundamental signal. We can use the relations shown in Figure 4 to estimate the loss in correlation coefficient, and hence, presuming the system temperature is not affected, the apparent loss in correlated flux density. These are shown in Table 1.

5 Experiment Setup

The primary goal of this experiment was to measure the coupling of the 3rd-order intermodulation response through observations of an astronomical source when the L-Band receivers on four VLA

Compression	Coupling	Loss	Loss (dB)
1%	-48.1 dB	0.99998	-.000067 dB
1 dB	-17.2 dB	0.98	-.082 dB
3 dB	-9.6 dB	0.90	-.45 dB
6 dB	-4.3 dB	0.73	-1.4 dB
10 dB	-2.0 dB	0.61	-2.1 dB

Table 1: Predicted loss of correlated flux density as a function of degree of compression.

antennas were driven into known degrees of compression by a high-power CW tone. This test attempted to simulate the strong radio interference we are likely to encounter with our future wideband EVLA receivers but under controlled test conditions and where the receivers are well characterized. The frequency of the interfering tone was selected so that the third-order intermodulation products from non-linear effects within the amplifiers would cause the strong OH maser line at 1665.5 MHz to be folded nearly on top of the HI line at 1420.5 MHz. Because the maser emission is very strong and narrow, and because its true strength could be easily measured independently, we could easily measure the phase an amplitude of the reflected response, and determine its coupling strength. Simultaneously, we could determine any other degradation of the true in-band emission (the continuum source and associated HI absorption).

Unlike the experiment described in EVLA Memo #79, which explored the effects of high-level, out-of-band RFI at C-Band, this test used an interfering tone lying well within the standard L-Band receiver frequency range used on the VLA. This presents technical challenges since it is the 3$^{\mathrm{rd}}$ order product that we are interested in and not the strong interfering tone itself. Consequently, the RFI signal which saturates the amplifiers in the receiver must be blocked from reaching, and saturating, any subsequent electronics in the signal chain.

For this experiment, the frequency of the RFI tone injected into the receiver was chosen to be 1545 MHz. The resulting 3$^{\mathrm{rd}}$order intermodulation product of interest to us was that appearing at a frequency given by

$$\nu = 2\nu_{RFI} - \nu_{OH} \tag{15}$$

which will create an image of the 1665 MHz OH line at a frequency of 1424 MHz – within 3 MHz of the HI line. Note that at the same time, the HI line will be reflected upwards to within 3 MHz of the OH line, due to the additional 3$^{\mathrm{rd}}$order product:

$$\nu = 2\nu_{RFI} - \nu_{HI} \tag{16}$$

which will create an image of the HI line at 1669 MHz. This latter reflection is of less interest to us because the strength of a typical HI line is much weaker than the OH maser line. Accordingly, given a suitable chosen astronomical source, the first intermodulation product will be much easier to detect and measure, and is the one we have used for this experiment.

There were a number of logistical constraints and system requirements that were met in order to carry out the L-Band compression test:

- Care was taken to ensure minimal disruption to the operation of any of the VLA L-Band front-ends and their associated LO & IF sub-systems. All modifications made on an antenna had to be low risk, easy to implement and quickly removable.

- As we planned to perform a phase-closure analysis of the interferometric data, four modified antennas were required. Since we didn't have access to enough individual laboratory frequency synthesizers to outfit each of the antennas with a strong 1545 MHz tone, the existing VLA system had to be adapted to provide the required CW tone. This was done by using one of the two L6 Synthesizer modules on each antenna.

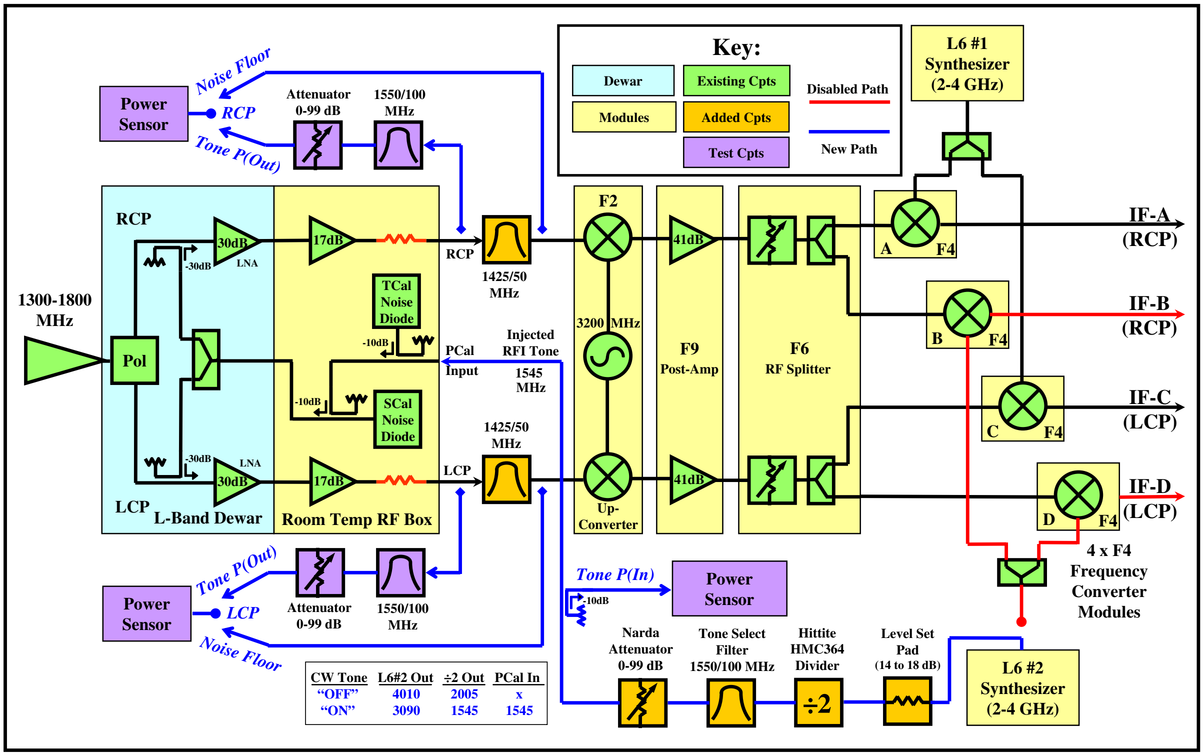

Figure 5: A block diagram showing the modifications to the VLA/VLBA L-band receivers utilized for this experiment. Block components are color coded to show the modifications made. The artificial tone was generated by one of the two L6 synthesizers, divided by two, and inserted into the signal path through the calibration couplers. Band limiting filters (in gold) were added to permit turning off the saturating tone by retuning the L6 synthesizer, and to prevent tone power from saturating subsequent stages of the electronics, as described in the text.

- To permit 'on-sky' comparisons, remote control of the CW tone was required in order to enable and disable the saturating signal. This was done using the VLA 'Observe' file to control the L6 synthesizer frequency settings so that a judiciously chosen filter would select or reject the interfering tone.

- To measure the coupling efficiency of the intermodulation product as a function of the degree of compression, the experiment was performed with three different levels of amplifier compression – nominally near 1 dB, 3 dB, and 6 dB, on both polarizations of the four modified antennas.

Figure 5 shows the block diagram of the L-Band compression experiment setup implemented on each antenna. The red traces indicate signal paths that were disabled or components that were removed. Blue traces and gold boxes show signal paths and components that have been added for the injection of the CW tone. Purple boxes portray components and test equipment used for characterizing the receiver's compression curve.

The saturating tone was provided by the second L6 synthesizer normally used to drive the B & D IF channels of the F4 frequency converter modules. The L6 synthesizer is restricted to a frequency range of 2 to 4 GHz with lock points roughly spaced every 20 MHz. The 1543.5 MHz tone needed to fold the 1666 MHz OH line exactly on top of 1421 MHz HI line thus lies well outside the L6's frequency range. However, the legal setting of 3090 MHz, when divided by two, provides an acceptable tone which would yield a separation of only 3 MHz between the HI and the folded OH line.

The divide-by-two function was provided by a Hittite HMC364 pre-scalar chip mounted on a microstrip circuit with SMA input and output connectors. Coaxial pads were used to attenuate the

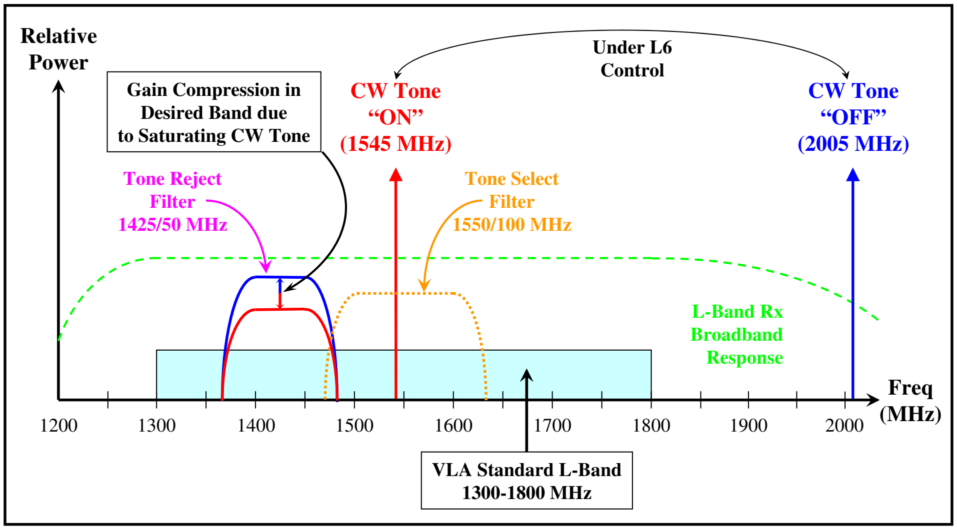

Figure 6: The simulated RFI tone was generated by dividing the 3090 MHz output of one of the L6 synthesizers by two. By using a 1550/100 MHz filter, the saturating tone was effectively turned off by retuning the L6 to 4010 MHz.

output level of the tone from the L6 which, when measured at the receiver, was found to lie in the 14 – 18 dBm range, depending on the antenna. The resulting 0 dBm signal was then fed to the Hittite board, which in turn delivered an output of about +7 dBm. The tone, with its frequency now divided by 2, went through a coaxial bandpass filter. This 1500-1600 MHz filter not only eliminated higher order harmonics generated in the divider but was effective in allowing the tone to be turned on or off remotely. When the L6 was set to 3090 MHz, the resulting divided tone at 1545 MHz, needed to generate the desired third-order intermodulation response, was passed through the filter. When the L6 was commanded to its highest allowed frequency of 4010 MHz, the resulting 2005 MHz signal was blocked by the filter, thus removing the saturating tone. The necessary commands were contained within the normal VLA 'Observe' file.

The tone select filter was followed by a 0-99 dB step attenuator which was used to set the power level of the saturating 1545 MHz tone. This provided a convenient way of mapping the compression curve of each receiver and allowed the CW tone power to be set in a repeatable fashion. A 10 dB coupler was added at the output of the step attenuator to allow the power level of the interfering tone, P_{in}, to be monitored. The RFI tone was added to the RF signal through the noise diode calibration path, which is accessible via the Phase Calibration Input port[4], which coupled the tone power to the noise calibration signal through a 10 dB coupler. The 1545 MHz tone then followed the standard noise calibration path, which includes a 3 dB splitter that feeds the tone to each polarization channel, and 30 dB calibration couplers which injected the tone onto the astronomical signal directly in front of the LNAs.

Figure 6 shows the details of the experiment on a frequency plot. The dashed green line outlines the broadband response of an L-Band receiver. The light blue shaded box shows the standard VLA 1300-1800 MHz observing band. When the L6 was set to 4010 MHz, the generated tone was at 2005 MHz (dark blue), and was rejected by the 1500-1600 MHz Tone Select Filter (shown in orange). Hence, at this setting, the saturating tone was 'Off'. When the L6 was set to 3090, the generated tone was at 1545 MHz (red), which passed through the Tone Select Filter, and into the receiver. Hence, at this setting the tone was 'On'. The resulting RF spectral power densities are schematically shown in blue for the uncompressed state, and in red for the compressed state.

[4]This input port is a VLBA feature which is not used on the VLA.

To ensure that only the FE was put into compression, and to prevent the powerful 1545 MHz tone from saturating downstream electronics, the tone power was prevented from propagating beyond the receiver by inserting a 1400-1450 MHz bandpass filters at the output of the receiver. This stripped away the 1545 MHz RFI tone but preserved the desired astronomical signal around the HI line frequency.

When considering the effects of compression in an amplifier, it is generally expected that the output power of the unit integrated over its total frequency range will remain constant. When a strong CW signal is fed into an amplifier, it will generate many harmonics and intermodulation products. As the input level is increased, more and higher level harmonics are generated (at least until the harmful damage point is reached) which will effectively reduce the gain seen at the fundamental frequency of the tone, as shown in Figures 3 and 4. This reduction is how the 1 dB compression point of an amplifier is both defined and determined. For low-level signals, the P_{out}/P_{in} ratio is, by definition, the gain of the amplifier. As the amplifier begins to saturate, an increase in the input level will no longer produce a proportional increase of the output signal. When the apparent gain drops by 1 dB, the measured output power level is thus defined as the amplifier's 1 dB compression specification (often denoted as the P_{1dB} point). The single-ended cryogenic low-noise amplifiers which we use in VLA L-Band receivers typically have a P_{1dB} of 0 dBm. Thus if the amplifier has a gain of 35 dB, when the input level is -34 dBm, its output will be at 0 dBm. That is, the amp will exhibit an effective gain of 34 dB when it is operated at its 1 dB compression point.

As will be discussed in detail later, the L-Band receivers were driven to their 1 dB compression point when the CW tone was typically at -15 dBm. It turned out that it was not the LNAs that are saturating in this experiment, but the first stage of post-amps following the LNAs. A back-of-the-envelope calculation of the power levels being experienced by the amplifiers in the signal path when the receiver is at its 1dB compression point is given in Table 2

L-Band Receiver Signal Path Location	Loss or Gain (dB)	Power Level (dBm)
Tone Power measured by Power Meter (P_{in})	–	-15
Input of Meter Coupler	+10	-5
Output of P_{cal} Coupler	-10	-15
Output of Cal Splitter	-3	-18
Output of Cal Coupler	-30	-48
Output Level of LNA ($P_{1dB} \sim 0$dBm)	+35	-13
Output Level of Post-Amp ($P_{1dB} \sim -5$ dBm)	+17	-4

Table 2: Estimates of typical power levels occuring in the L-Band compression test.

At the 1dB compression point, the input power at the LNAs was around -48 dBm while the resulting output was about -13 dBm, which is more than 10 dB below the 1 dB compression point of 0 dBm for the low noise amplifier. The post-amp output, however, has a measured P_{1dB} of -5 dBm. When the output is at -4 dBm, the post-amp will be at its 1 dB compression point.

A picture of the setup used in the lab is shown in Figure 7. This shows the Hittite divide-by-2 board along with the use of the P_{cal} input port. Figure 8 shows the actual test setup as implemented on Antenna 23. Of note is the step attenuator along with the 1500-1600 MHz filter and 10 dB test coupler. The C-clamp that is used to mount the test tone components may seem rather jury-rigged but in fact provides a simple yet functional way to execute the required modifications.

One advantage of using the P_{cal} path for injecting the 1545 MHz RFI tone was that the T_{cal} feature was preserved. Unlike the earlier C-Band compression test described in EVLA Memo #79, this meant that the switched noise power signal remained available for calibrating each L-Band system during the experiment. One might assume that in a saturated receiver, the change in power between the Cal Off and Cal On (typically set to about 10% of T_{sys}, which corresponds to several tenths of a dB increase

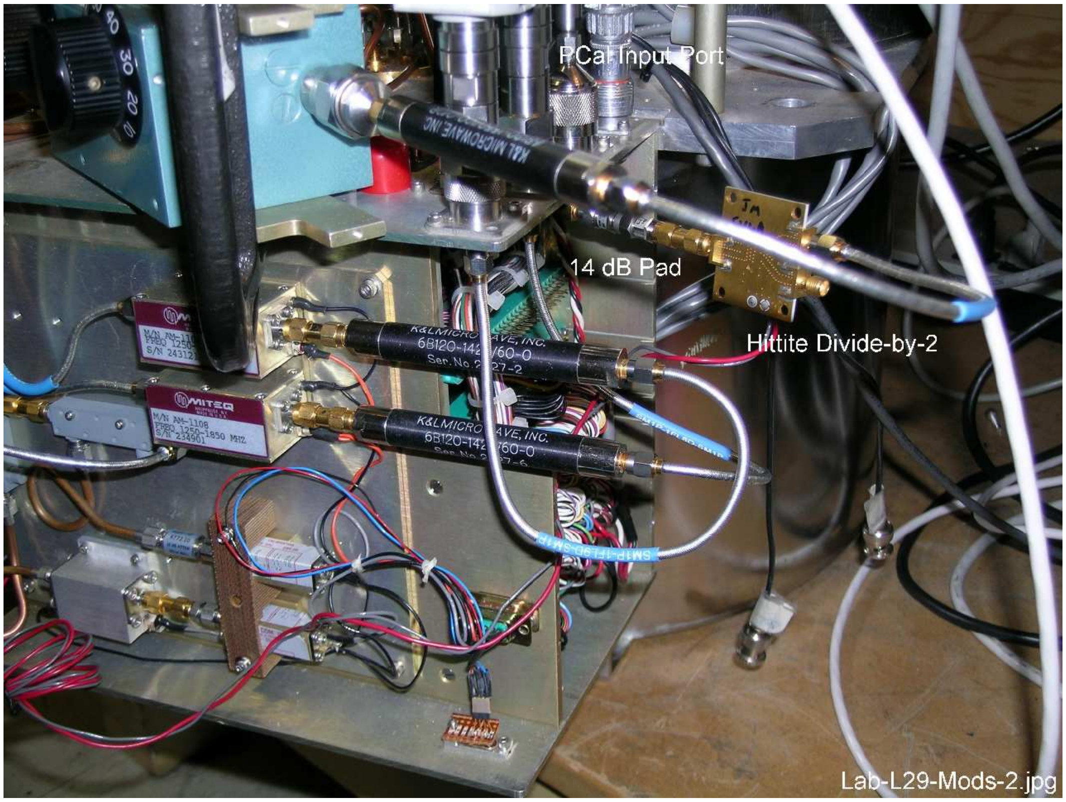

Figure 7: Picture of the lab test setup, showing the Hittite divide-by-2 circuit board

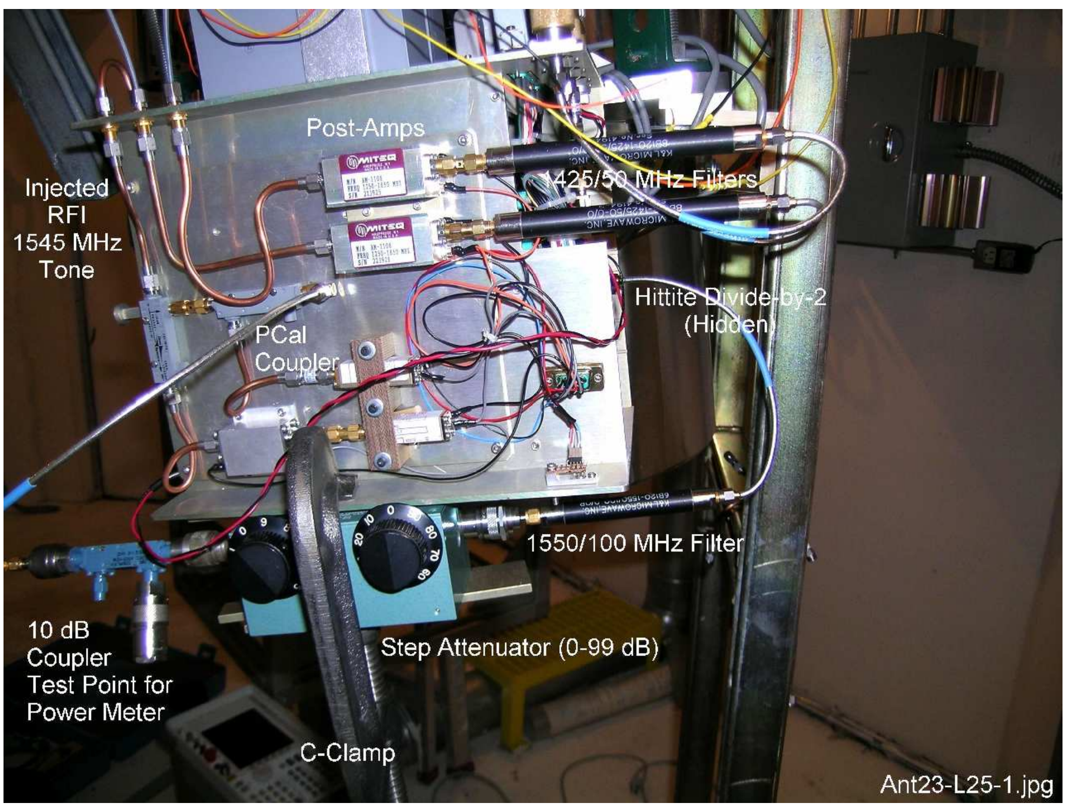

Figure 8: Picture of the modifications used on Antenna 23

in output power on cold sky) would be affected. Tests in the lab at C-Band in 2004 showed this is definitely not the case at the levels of saturation employed in our tests. The difference in switched power was measured on a VLBA C-Band receiver with increasingly higher LNA saturation levels (up to a compression level approaching 8 dB) from an out-of-band CW tone. No noticeable change was seen in the T_{cal} power delta, at least down to the 0.01 dB measurement level (i.e., 0.2%). While the absolute power of the amplifier changed as a result of the gain being compressed, the T_{cal} change in power remained constant. This result would not be the case if, for example, the amplifier in question was being saturated by noise power from cold sky (rather than a CW tone). In this scenario, the T_{cal} switched power change would undoubtedly be 'squished'.

6 Laboratory Tests

Figure 9 shows the results of a verification of the experiment concept that was carried out with a spare L-Band receiver (L#29) in a lab setting. Two Agilent 83630B synthesizers were used to simulate the 1421 MHz (HI) and 1666 MHz (OH) spectral lines. The signals were combined with a splitter and fed directly into the receiver through the quadridge OMT test fixture normally used to provide hot/cold noise standards. An 83623L synthesizer provided the 3090 MHz RFI tone to the Hittite divide-by-2 which was injected into the receiver using the noise calibration path, although in this case, through the solar cal path rather than the P_{cal} coupler as described above, thus providing 10 dB higher tone levels for a given attenuator setting than were seen in the antenna tests performed later. The output of the receiver was measured with an Agilent 8563E spectrum analyzer which was configured to duplicate what the VLA correlator would see in spectral line mode with a 6.25 MHz bandwidth centered at 1422.5 MHz or at 1667.5 MHz. The spectrum analyzer provides 601 spectral channels - many more than the correlator has - with a 10.4 KHz spacing and a resolution bandwidth of 100 KHz.

Figure 9 shows the spectra at both the 1422.5 and 1667.5 MHz center frequencies for the LCP and RCP channels for various attenuator settings. The larger the attenuator value, the less power in the 1545 MHz RFI tone. The traces at the top show the situation before any compression occurs. One can see the simulated HI and OH lines at 1421 and 1666 MHz respectively in both channels. The strength of these lines was set to be about 10 dB above the noise floor, which required a power level at the input of the receiver of -101 dBm. As the power of the RFI tone increases, the third-order products begin to appear. The OH line is folded down to 1424 MHz (i.e., $2 \times 1545 - 1666 = 1424$ MHz) while the HI line is folded up to 1669 MHz (i.e., $2 \times 1545 - 1421 = 1669$ MHz). Note how the strength of both the HI and OH lines drop, as does the noise floor level (at least for a while), as the aliased lines grow in amplitude. It is also obvious that the LCP channel compresses before the RCP side does.

An Excel spreadsheet was used to calculate the average power in the noise floor around each of the primary and folded lines, as well as to determine the maximum power of each of the four lines (within a (0.25 MHz window) for the various attenuator settings. Figure 10 shows a plot of the noise floor compression ratio (green traces), the 1421 & 1666 MHz line compression (dark and light blue traces) and the 1444 & 1669 intermod line expansion (red & orange traces) with the LCP and RCP curves shown separately.

On the RCP side, the noise floor drops steadily as the power of the injected RFI tone increases. The amplitudes of the 1421 and 1666 MHz simulated spectral lines also fall while the 1424 and 1669 MHz intermod lines grow and eventually reach a plateau. All of these effects are expected on the basis of the analysis given in Sections 2, 3, and 4. However, on the LCP side, the noise floor initially drops as tone power increases, but then begins to increase as the tone power reaches its highest levels. In the end, the noise floor had more average power than it started with. The 1421 and 1666 MHz lines drop, and unlike the RCP channel, leveled off. The 1424 & 1669 MHz itermod lines stop growing and eventually plateau, much as they did on the RCP side.

The rising of the noise floor at high compression levels on the LCP channel came as a surprise. But

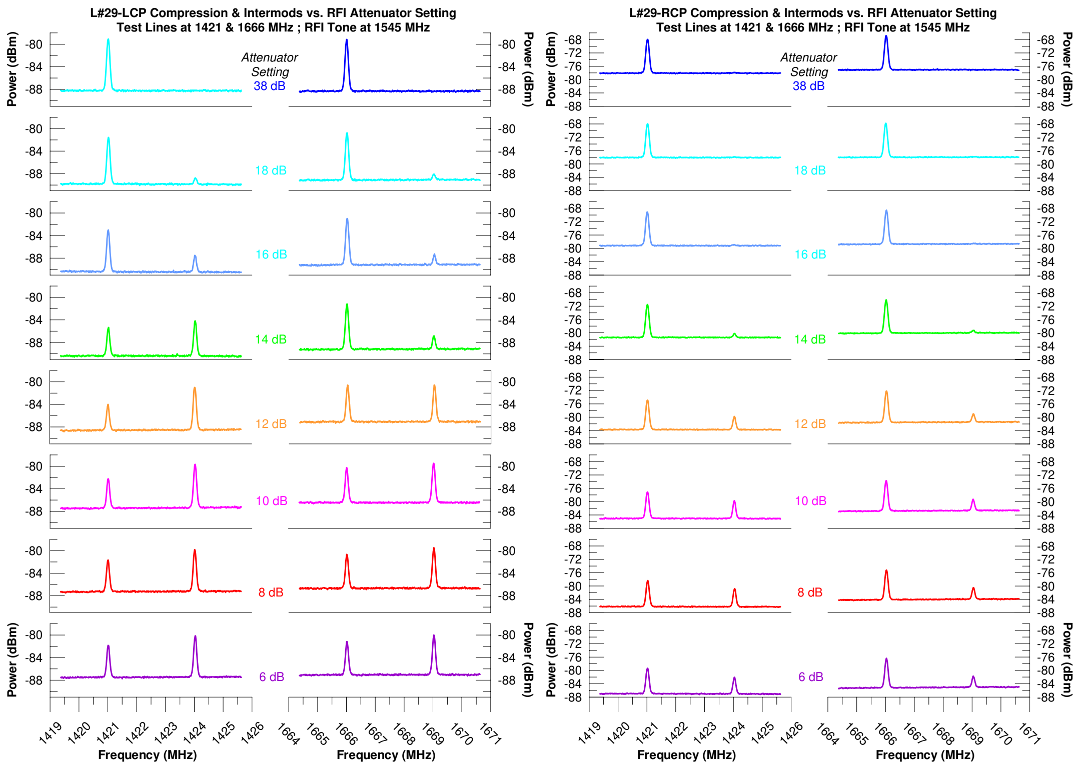

Figure 9: Simulated lab tests on L#29 in LCP (left) and RCP (right). The 3rd-order intermodulation products at 1424 and 1669 MHz, due to coupling of the 1421 and 1666 MHz tones with the 1545 RFI tone, are seen to rise as the power in the 1545 MHz tone is increased, from top to bottom.

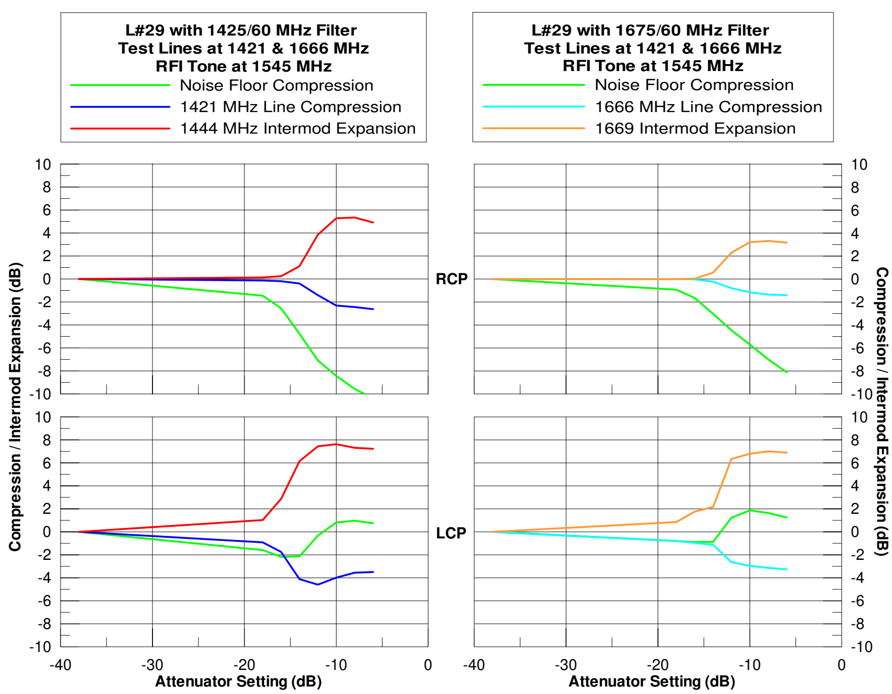

Figure 10: Noise Floor and Line Compression analysis of the L#29 lab test.

as this polarization saturated much earlier than the RCP channel, this seems to suggest that when an amp is driven heavily into compression, other effects, not included in our simple model, begin to dominate. At low levels of saturation, one should be able to measure the level of compression by looking at how much the power level of the noise floor drops. This was the method that was successfully used in the earlier C-Band compression experiment (described in EVLA Memo 79) to determine the LNA saturation level. This method makes it quick and easy to characterize the compression curve but would obviously not be accurate enough to quantify the effects of the strong simulated RFI tones planned in the L-Band experiment.

To explore the increase in spectral power at high levels of compression, further lab tests were carried out on the L#29 front-end using a spectrum analyzer which looked at the output of the receiver across a wider frequency range covering 1405-1705 MHz with a resolution bandwidth of 1 MHz. This range not only includes the 1421 and 1666 MHz bands as well at that of the RFI tone, which for this experiment was set to 1555 MHz (this frequency provides a wider separation between the fundamental and aliased lines). The measured noise floors for LCP and RCP are shown in Figures 11 and 12.

The yellow bands along the X-axis show the frequency coverage of the 50 and 60 MHz wide filters used at 1425 and 1675 MHz during the experiment. The dark blue traces show the noise floor with no compression. As the strength of the RFI tone is increased, the noise floor level initially drops slowly but steadily. But beyond a certain tone level (shown by the the green trace), the noise floor begins to show an undulating pattern, with a dip close in to the tone that rises quickly on both sides. It then tapers off slowly but with the noise floor elevated in the 1421 and 1666 MHz bands above the non-compressed level.

This demonstrates how the noise floor varies with RFI tone strength, although it does not explain why. The reason is likely due to the production and interaction of higher order products arising in the amplifiers when in a highly compressed state. For this experiment on the VLA, the main consequence is that the compression curve of each receiver can't be done by simply measuring the drop in the noise floor as the tone strength increases. The system must be characterized by determining the direct compression of the RFI tone by measuring the P_{in} and P_{out} levels, plotting the P_{out}/P_{in} ratio and determining the level of compression from a normalized power curve.

7 Preliminary Antenna Tests

On December 7, 2005 a 'dry run' was performed on Antenna 23 to check out the viability of using the L6#2 module as the RFI tone generator. Experience was also gained in modifying the L-Band receiver in situ. All worked out well, although this exercise provided the first evidence that characterizing the receiver's compression curve using the noise floor technique was not going to work (note that all the earlier lab tests were done using the RCP side of L#29 which, as luck would have it, had a very benign and abnormally 'normal' saturation curve).

Figure 13 shows the noise floor compression curves found for two polarization channels of the L#25 receiver on Antenna 23.

A filter bank was used that allowed the spectrum analyzer to measure both the 1421 and 1666 MHz regions while rejecting most of the 1545 MHz tone (this was done to ensure the mixers in the analyzer aren't themselves overdriven). The filter bank consisted of a 2-way splitter with one arm feeding a 1400-1450 filter and the other a 1650-1700 MHz filter. The signals were then recombined in a second 2-way splitter. The spectra of the noise floor bandpasses illustrated in the upper plots shows that the non-compressed level (the blue trace) drops steadily down with decreasing attenuation (i.e., increasing 1545 MHz tone power) until about 8 dB (the green trace), at which the noise floor level begins to grow and, at maximum tone strength (the red trace), it is much higher than the non-compressed level.

An Excel spreadsheet was used to average the noise power in each of the 1421 & 1666 MHz filter bank bands and the resulting compression curves are shown in the bottom plots in Figure 13. The

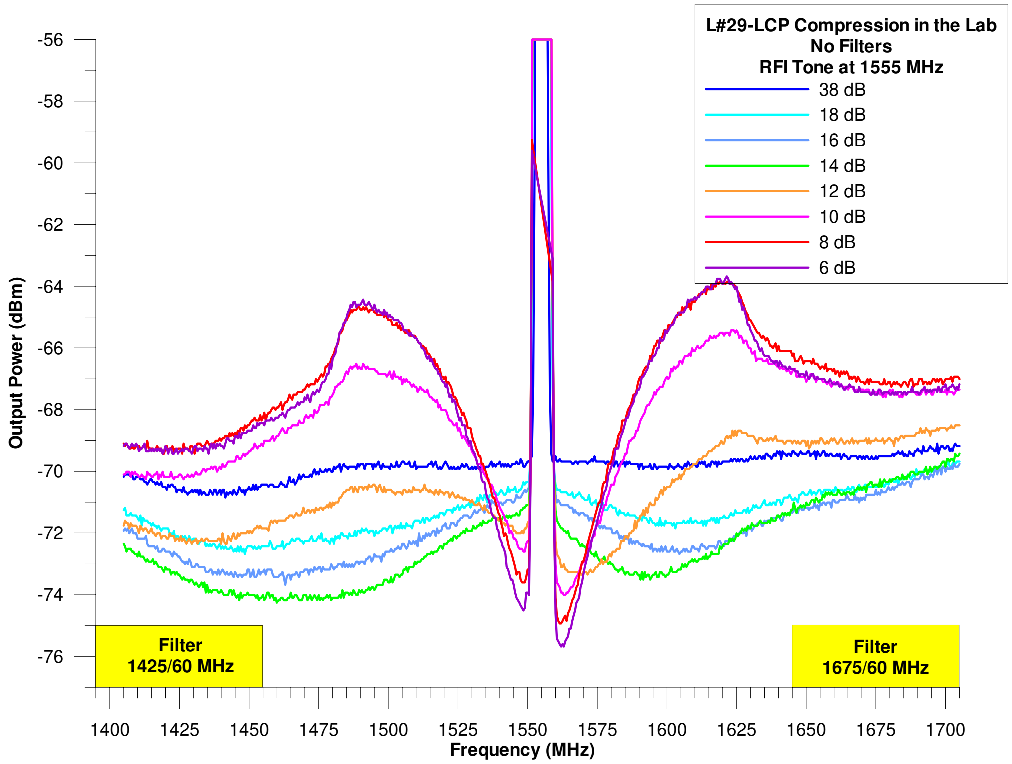

Figure 11: Frequency response of noise floor compression test on L#29-LCP

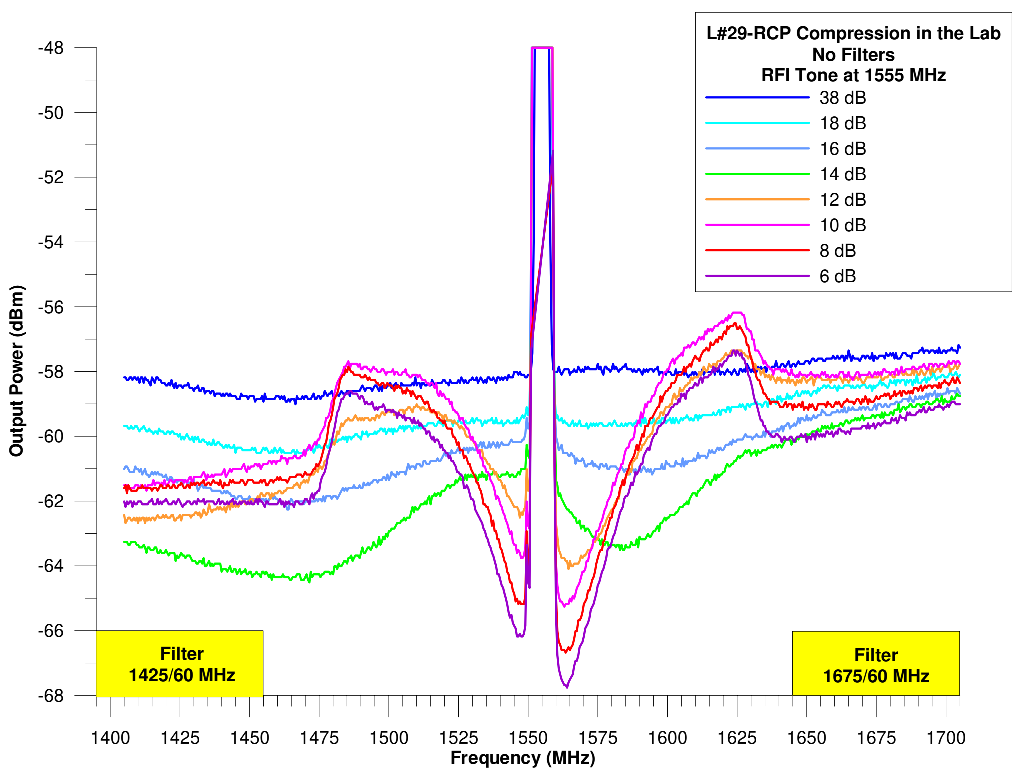

Figure 12: Frequency response of noise floor compression test on L#29-RCP

noise floor at 1421 MHz does compress for a while as the tone strength is increased but the noise floor at 1666 MHz hardly compresses at all and eventually starts to shoot up quickly. This effect is

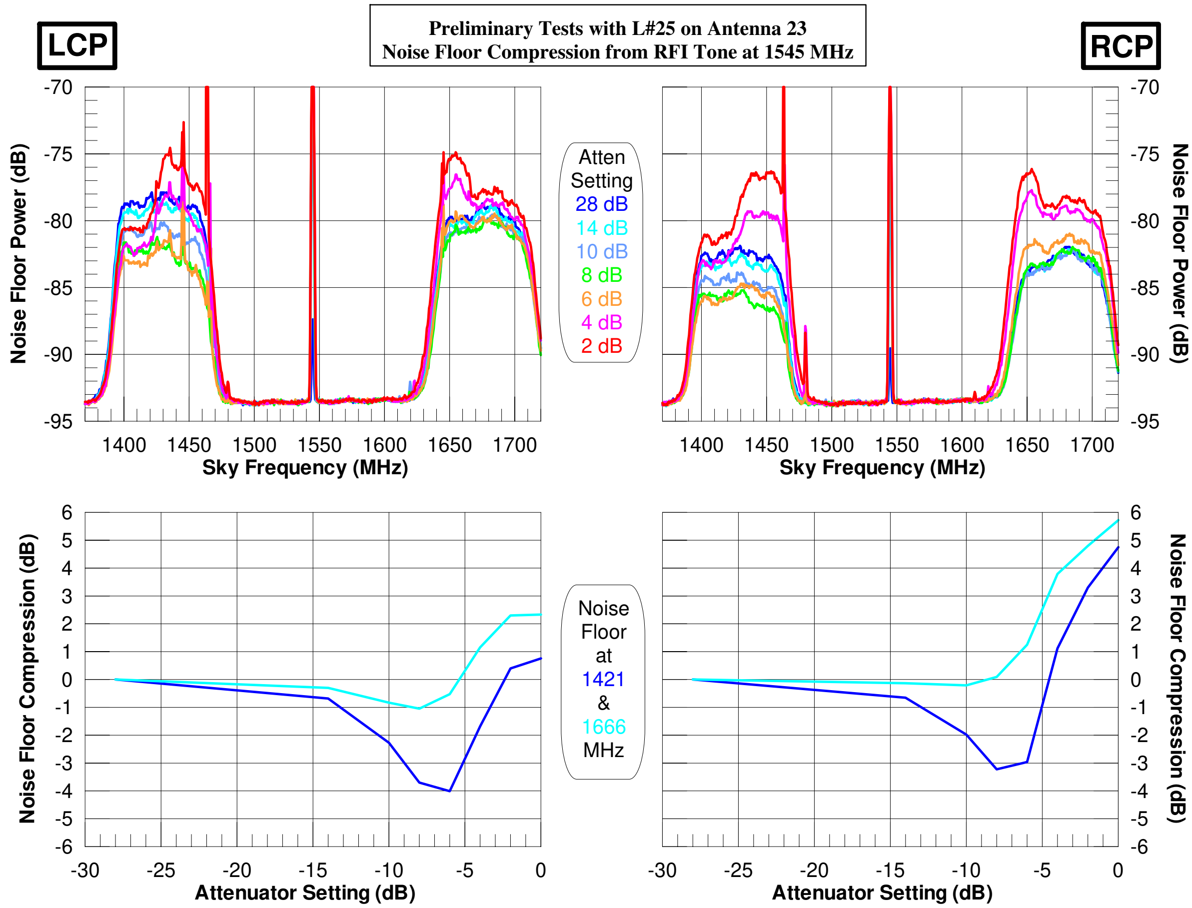

Figure 13: Measurements of the noise floor compression of L#25 on antenna 23.

spectacularly prevalent on the RCP side.

The discovery of this effect led to a much more comprehensive set of lab tests and the subsequent development of an Excel spreadsheet that would allow us to enter both direct tone compression measurements and noise floor data taken with a power meter test setup and essentially generate the compression curves in real-time.

8 Preparations for the On-Sky Compression Test

The on-sky L-Band compression observations were performed on January 3 & 4, 2006, when the array was in its D-configuration. This configuration was necessary not only because access is easier to the four modified antennas, but also because the shorter baselines result in lower fringe rates which minimized the effects of fringe washing in the correlated baseline visibilities.

Four VLA antennas – #9 at N8, #12 at E6, #23 at W6, and #27 at N2 – were modified for the experiment by a single engineering team on January 3. These modifications took about 30 minutes each. An additional 60 minutes was required to characterize the compression curve for each receiver. This procedure involved measuring the direct P_{out}/P_{in} ratio of the 1545 MHz tone as well as the variation of the noise floor power level in the 1400-1450 MHz filter band for a wide range of RFI tone attenuator settings on both polarization channels. An Agilent E4419B Power Meter with E4113A (0.1-18 GHz) detectors was used and the data were entered into an Excel spreadsheet which allowed near real-time mapping of the compression curve.

In order to minimize any non-linear effects coming from the power detectors, a second step attenuator was used to adjust the strength of P_{out} so that it was nearly identical to the P_{in} level. Since a true amplifier gain curve would require the detectors to be linear over many tens of dB, this scheme results in both P_{in} and P_{out} lying at the exact same operating point of the detectors so that the differential non-linearity is negligible. A 1500-1600 MHz filter is placed in front of the detector to eliminate most of the noise power coming from the receiver. This ensures that we are predominately measuring the amplified 1545 MHz tone.

Spectrum analyzer measurements were performed during maintenance time on the morning of January 4 while unrelated software problems were sorted out. These measurements are shown in Figure 14 and 15 individually for each antenna and for both the LCP and RCP channels. Each receiver channel had a filter centered at 1425 MHz to pass the astronomically interesting band but reject the tone at 1545 MHz so it doesn't lead to unwanted 3$^{\text{rd}}$order products in the down conversion chain that follows. Note that half of the filters are 8-pole 50 MHz bandwidth units while the rest are 6-pole 60 MHz filters. A search of the lab turned up fourteen of the 1425 MHz filters, all of which were checked on a network analyzer in order to pick the eight that had the best rejection at 1545 MHz. All of the selected filters attenuate the RFI tone by more than 70 dB. Note that Figures 14 and 15 suggest that only three of the 12 channels have noise floors that compress steadily with increasing tone power. The receiver outputs associated with RCP on Antenna 9 and both channels on Antenna 27 have the lowest noise floor level when the tone is at its strongest (i.e., the red traces are lower than any of the other traces). Antenna 9-RCP is the only channel that has no anomalous bump appearing in the passband.

Once the receiver was modified and the compression curve measurements were taken, the attenuators were adjusted to put the RCP channel of each receiver as close to the 1 dB compression point as possible. Since the amplifiers in the RCP & LCP paths have slightly different compression performance and since there is no provision to independently adjust the balance between the two noise calibration channels, it was impossible to produce exactly equal compression levels on both sides, so we had to accept whatever compression level the LCP side delivered.

Table 3 shows the compression levels achieved on the four antennas during the experiment. For the receivers in the 1 dB compression test, the RCP channel compressions ranged from 0.8–1.2 dB. The LCP channel exhibited higher compression levels in the 1.5 to 2.4 dB range. For the second test, which attempted to compress the RCP side to the 3 dB point, the RCP side compressions varied from 2.6–2.9 dB while the LCP side again yielded higher values in the 3.5–4.0 dB range. In the third and final test, which put the receivers into the maximum amount of saturation as the test setup would allow (i.e., the attenuators were set to 0 dB), the RCP channels ended up with compression varying between 4.4 to 7.2 dB while the LCP presented compression levels of 5.3–8.2 dB.

Antenna & Receiver		1 dB Compression Test			3 dB Compression Test			Max Compression Test		
		Att Setting (dB)	Tone Compression (dB)		Att Setting (dB)	Tone Compression (dB)		Att Setting (dB)	Tone Compression (dB)	
Ant	L#		RCP	LCP		RCP	LCP		RCP	LCP
9	35	9	0.83	1.64	4	2.81	3.59	0	6.19	7.26
12	23	6	0.99	2.21	2	2.89	3.61	0	4.39	5.28
23	25	9	1.12	2.40	5	2.77	3.50	0	7.24	8.25
27	10	5	1.16	1.48	2	2.58	4.04	0	3.91	5.69

Table 3: Measured L-Band 1545 MHz tone compression levels used in the experiment.

Figure 16 shows each of the receivers' output power for increasing power levels of the tone. The P_{out} versus P_{in} plot typically shows a sort of S-curve. The rollovers at the top right of the curves are from the amplifiers going into compression. At the low end, the power in the injected tone is not

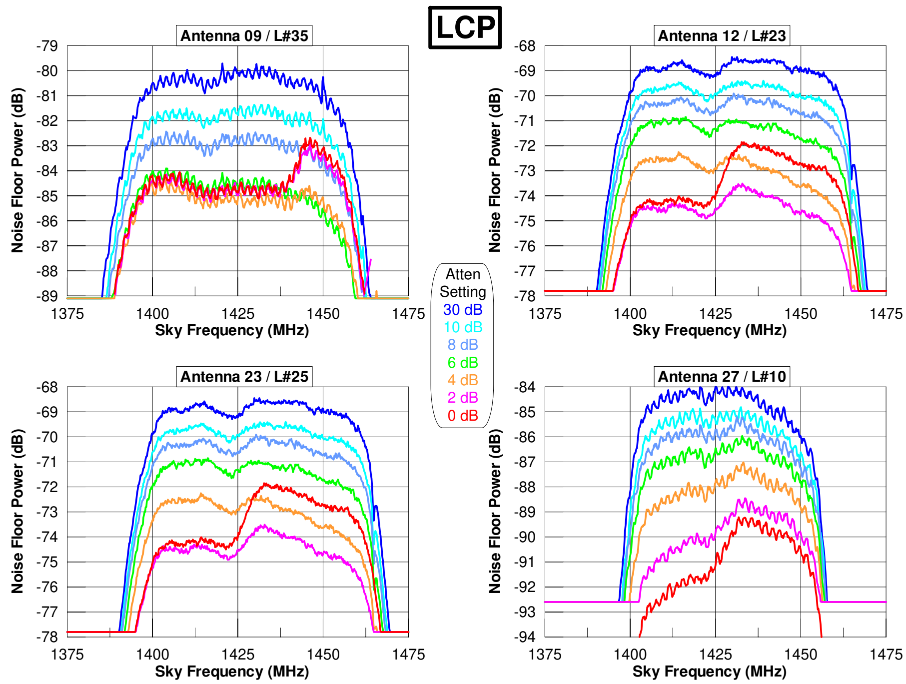

Figure 14: Noise floor power in the 1425 MHz filter versus attenuator setting for LCP channels

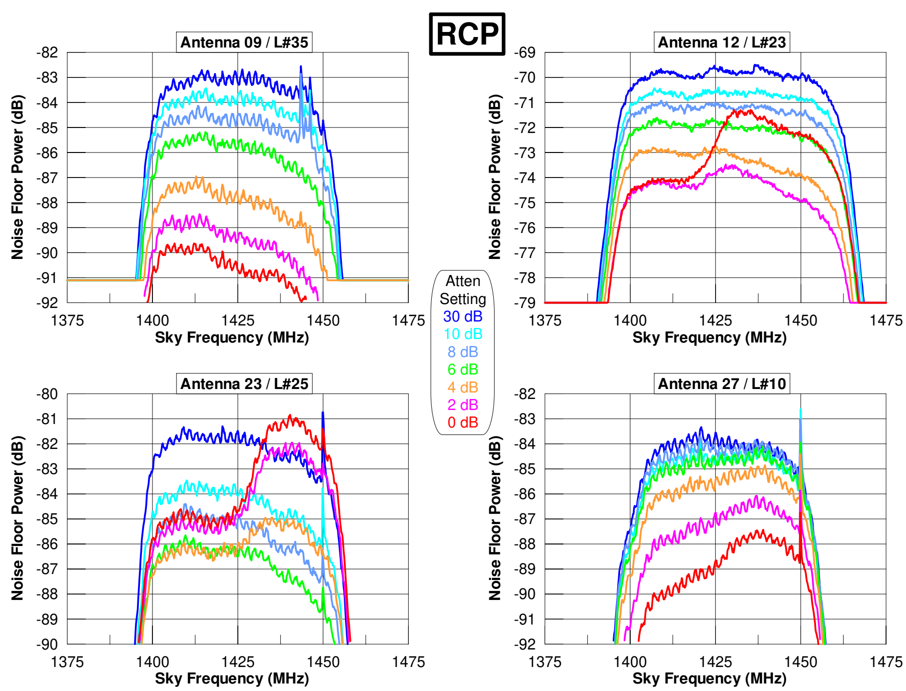

Figure 15: Noise floor power in the 1425 MHz filter versus attenuator setting for RCP channels

much different than the amplified noise power from the amplifier itself. Consequently P_{out} is essentially dominated by the noise figure of the amplifier chain. Thus the S-curve essentially shows the dynamic range of the system. The curves look well behaved with the exception of the RCP side on Antenna 12, whose clearly discrepant point we ascribe to a a typographical error during the recording of the data. It does not affect the more important compression information.

The apparent loss of gain due to harmonic power conversion, for both the saturating tone, and

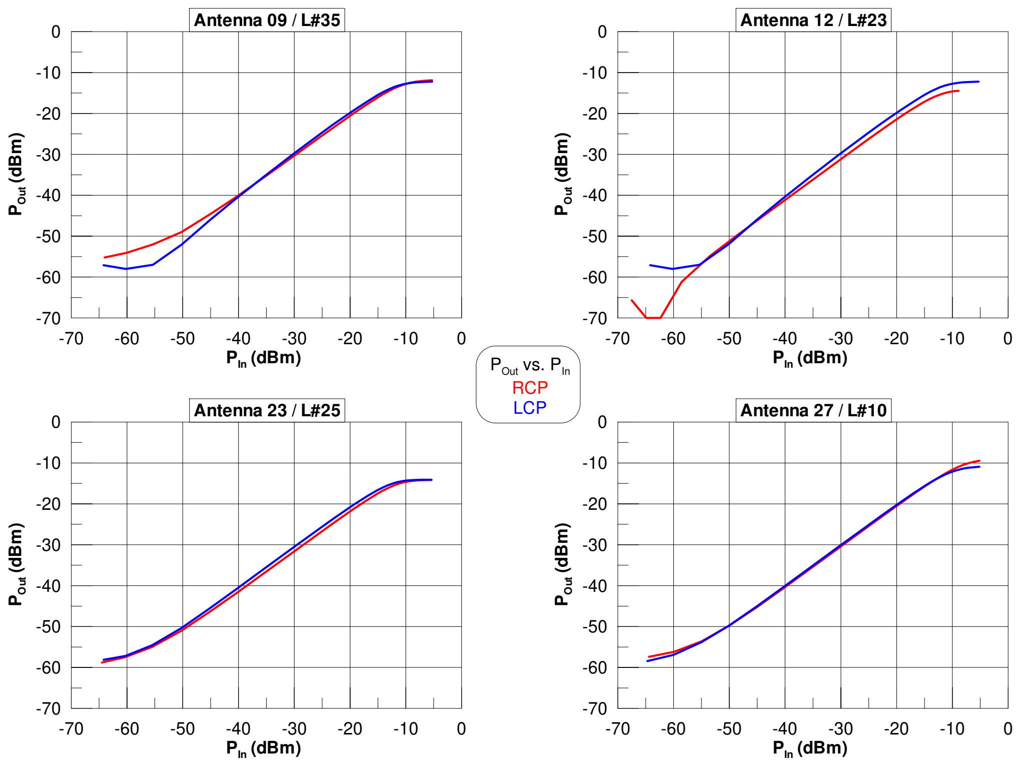

Figure 16: Measured P_{out} vs. P_{in} for the 1545 MHz tone.

the (astronomical) broad band noise are shown in Figure 17. We see here that for all amplifiers, the loss in noise power due to saturation exceeds that in tone power. This is expected on the basis of the saturation model described earlier. For small degrees in compression, the model predicts the noise power loss, in dB, to be double that in tone power. Inspection of the plots shows generally good agreement with this prediction.

In general, the curves all show similar behavior, with the anomalous turn up in compression level occurring at attenuator settings less than about 4 dB. However, as suggested earlier in Figure 15, only the RCP channel on Antenna 9 and both sides of Antenna 27 show no evidence of rolling up. If one assumes that systems that have higher compression points, and are thus not as susceptible to a saturating tone, would be less likely to exhibit this behavior, then Antenna 27 might prove this theory is valid since the Max Compression Test data in Table 3 shows it is the receiver with the least amount of compression. However, the RCP side of Antenna does not look particularly special. It shows the 4th worst compression level under the Max Compression condition yet its noise floor compression curve looks better than that on Antenna 27.

Our amplifier compression model predicts the astronomical noise loss to be double that of the tone compression loss, for small degrees of compression. The amplifier compression measurements shown in Figure 17 can be recast in a form suitable to judge the validity of our compression model. The results are shown in Figure 18.

All amplifiers show a linear relation between noise and tone attenuation for tone compressions less than 2 dB. Four of these have a slope near the expected value of 2, while the other four have slopes notably less than this. An explanation for the lesser slope might be that the actual amplifier gain for the tone at 1545 MHz is different than that for the noise power at 1425 MHz – a degree of freedom not incorporated in our model, which has assumed a single gain value for all frequencies.

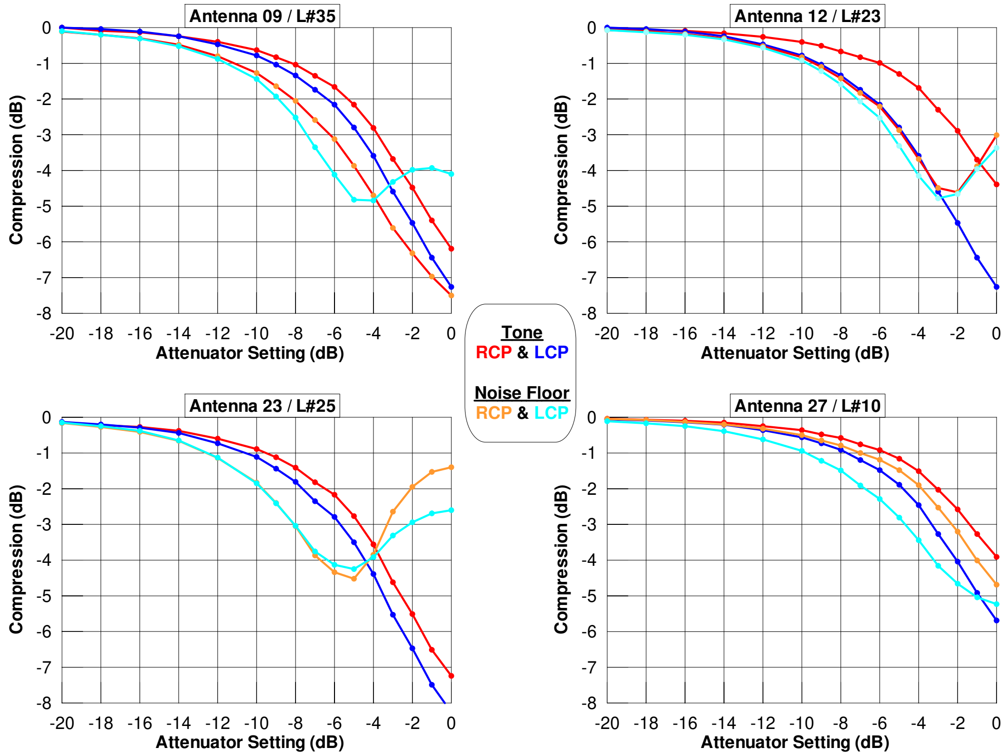

Figure 17: The measured 1545 MHz tone and noise compressive curves, showing the loss in power of both the saturating tone and the astronomical noise power due to conversion to higher harmonics.

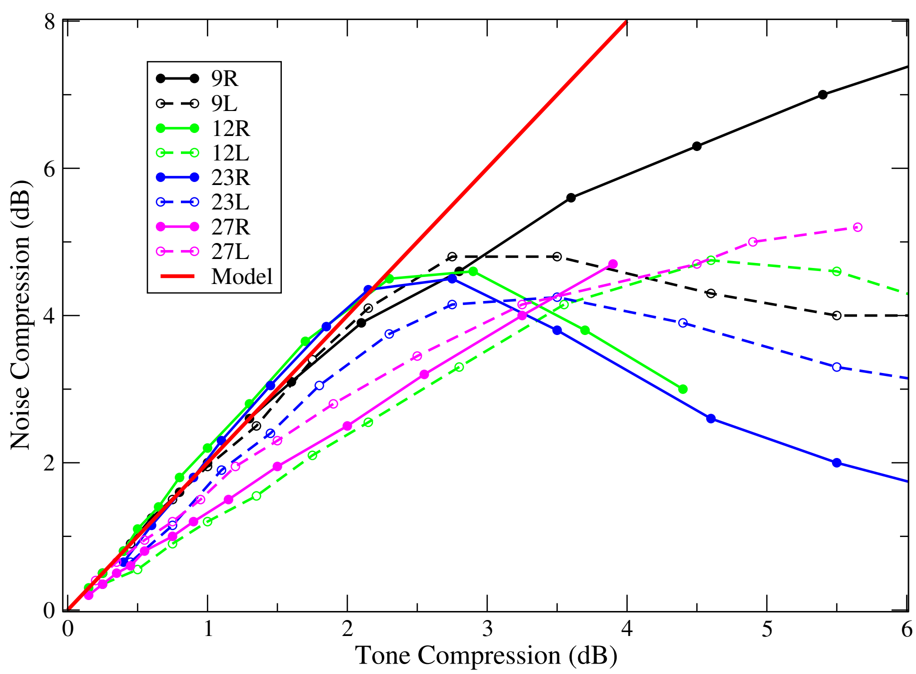

Figure 18: The loss in noise power plotted against the loss in tone power (compression). The straight red line shows the prediction of our amplifier compression model. The large deviations from linearity for most amplifiers at high compression is a result of the additional noise power from an unmodelled process.

9 The Astronomical Results

The on-sky observations were taken on the afternoon of 4 January, 2006. We observed the source W3OH, which includes both a strong (200 Jy) maser, and a strong (total flux of about 30Jy) continuum source, allowing easy detection of HI in absorption. Observations were taken at the center frequencies and resolution bandwidths given in Table 4.

Observational Setup						
Freq. MHz	BW MHz	$\Delta\nu$ kHz	N_{ch}	Mode	Avg. sec	Comments
1422.493	6.25	48.8	128	1A	6.7	Both HI and aliased OH in RCP
1422.493	6.25	48.8	128	1C	6.7	Both HI and aliased OH in LCP
1420.512	1.56	6.10	256	2AC	13.3	Centered on HI absorption
1424.474	1.56	6.10	256	2AC	13.3	Centered on OH alias

Table 4: The observational setup

The observations were made both with the tone on and off, with the tone power set to put the amplifiers into approximately 1, 3 and 6 dB compression, as described in the preceding section. The observations at each of these levels took 48 minutes, including a calibration scan on 3C48 taken with the tone 'off'. Following the '1 dB' and '3 dB' observation blocks, the tone power attenuators were adjusted to the values required for the next compression level, requiring approximately 20 minutes each. Following the '6 dB' observation block, all antennas were returned to their normal operating state.

The compression levels were measured in the field on the preceding day, and are given in Table 3.

9.1 Data Calibration

The data were calibrated using standard procedures. Observations of the calibrator source 3C48 (with the saturating tone off) were used to determine the spectral bandpass, and to set the correct gain. 3C48 is unresolved to the array resolution of $\sim$ 60 arcseconds, so no corrections for resolution were required. Data quality was excellent, and virtually no flagging of bad data was necessary.

Maps of the source emission, with and without the maser emission, are shown in Fig. 19. The left panel of Figure 19 shows the continuum emission at 1424.46 MHz, the right panel the emission at the maser frequency of 1655.591 MHz. At this frequency the continuum emission is unchanged (the notable emission at the north end of the 1424.5 MHz image is absent in the 1665.6 MHz image due to attenuation by the primary beam), and the unresolved maser is the dominant structure. The maser emission is circularly polarized, with I = 172 Jy and V = -28 Jy.

The OH emission spectra at both resolutions and polarizations are shown in Fig. 20.

The averaging time was set to the minimum time permitted by the correlator for the mode required – 6.66 seconds for the 1A and 1C observations, and 13.3 seconds for the 2AC observations. At the time of this experiment, we were unaware that the phase rate of the aliased response could be faster than that of the 'natural' rate. In fact, as explained below, for this experiment the rate is about 1.5 times that of the natural rate, resulting in a significant amplitude attenuation of the aliased signal, especially for the 2AC mode, where the integration time was 13.3 seconds.

10 Results

All anticipated consequences of the amplifier compression were easily detected. The observations clearly showed the expected aliased OH signal, with a strength that is a sharp function of saturation level.

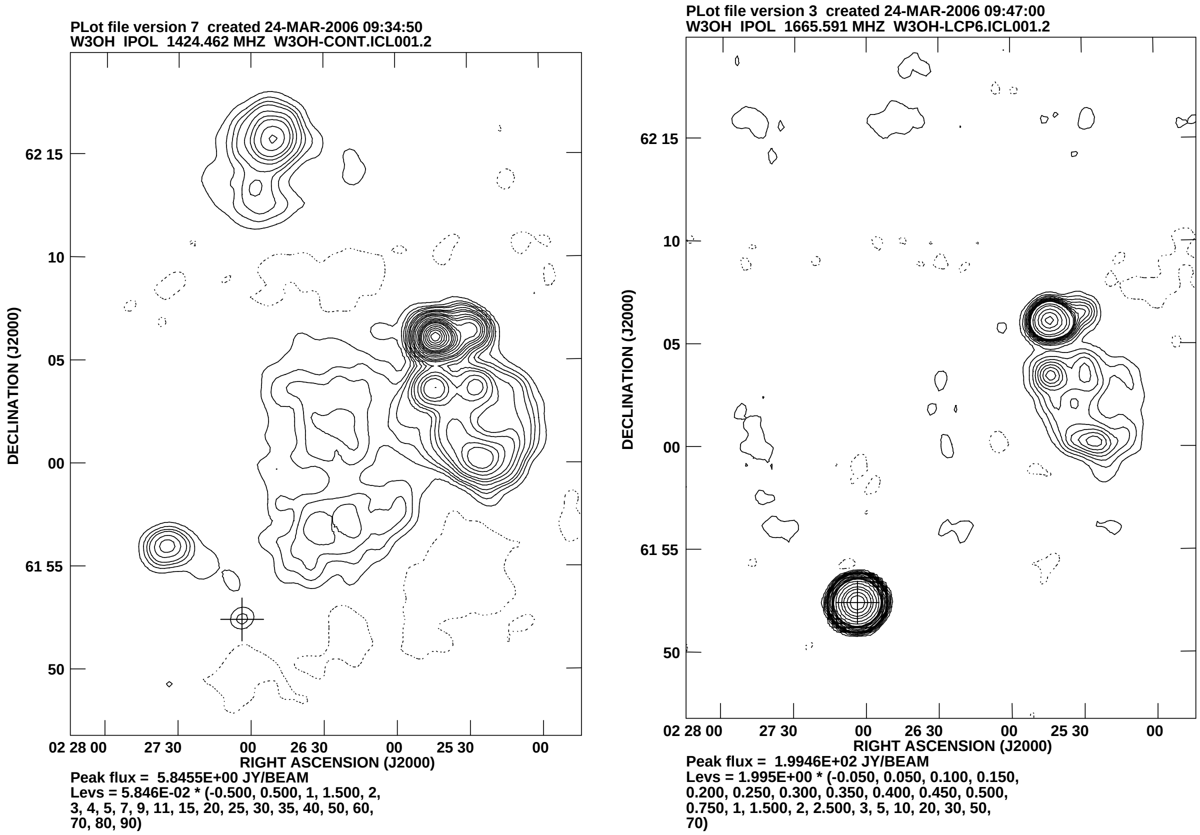

Figure 19: (Left Panel) A map of the target source W3OH showing the continuum emission at 1424.5 MHz, with 60 arcsecond resolution. The location of the maser source is shown with the 'plus' in the lower left. (Right Panel) A map of the LCP emission at the maser frequency of 1665.591 MHz, and a channel width of 6 kHz. The continuum emission is unchanged, while the maser emission is unresolved in both space and frequency.

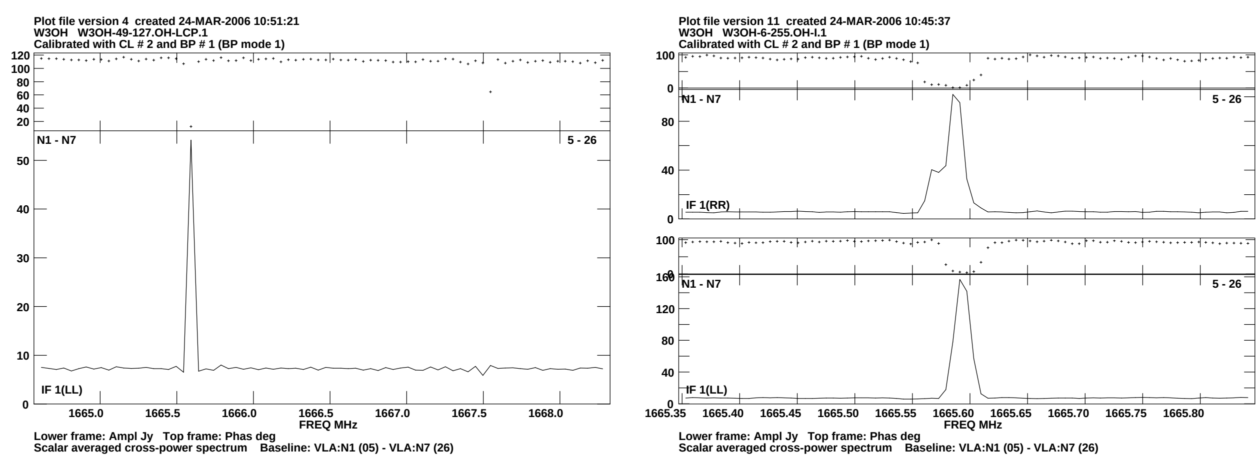

Figure 20: (Left) The OH emission spectrum at a resolution of 49 kHz. The strong 1665 MHz emission is seen prominently on the left side. The much weaker 1667 MHz feature is visible – more prominently in phase than amplitude – on the right side of the spectrum. The continuum emission causes the offset in amplitude, and the non-zero phase, of the remaining channels. (Right) The RCP (top) and LCP (bottom) emission profiles at a resolution of 6 kHz.

The phase of the aliased signal rotated as expected, and the visibility amplitudes of the non-aliased signal were reduced.

We discuss these in turn.

10.1 Coupling of Out-Of-Band Emission by Intermodulation

The presence of the strong RFI tone acts like a local oscillator, and reflects out-of-band emission into the observing band. The effect is shown in Figure 21. The left panel shows a piece of the normal

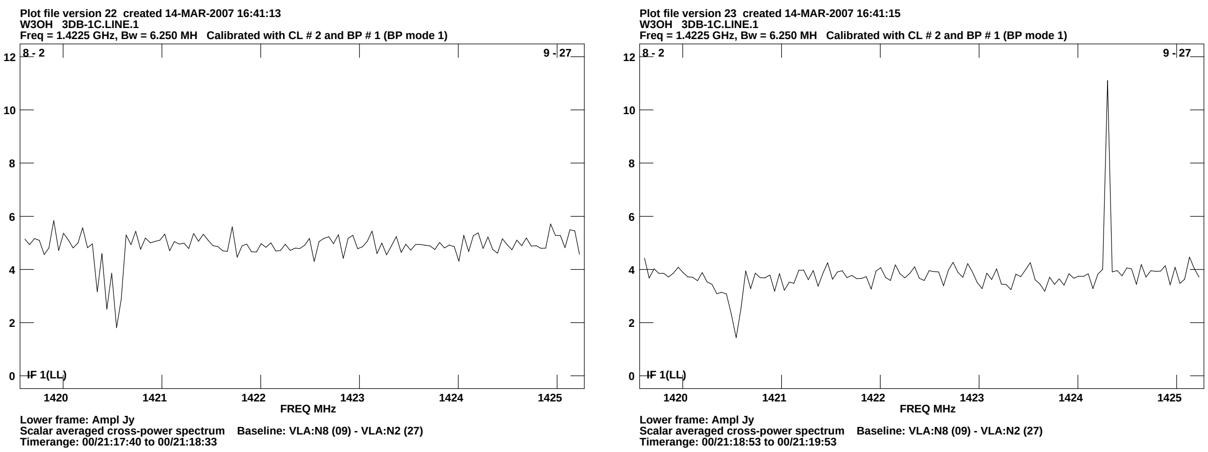

Figure 21: (Left) A single-baseline spectrum in LCP between antennas 9 and 27, with the tone off, showing the HI absorption and continuum. (Right) The same baseline and polarization with the tone on. A reflection of the OH emission appears at 1424.5 MHz, while the flux density of the continuum and HI emission declines.

spectrum, with the HI absorption clearly visible on the left side. The right panel shows the emission spectrum when the tone has put the amplifier into 3dB compression. Two changes are immediately apparent: The appearance of an image of the OH emission at a frequency of 1424.25 MHz, and an overall reduction in the visibility amplitudes of the in-band emission.

The coupling efficiency of the reflected OH emission can be directly measured from the strength of OH line, correcting for the loss due to coherent phase winding, and knowledge of the true emission strength. The integration period of 6.6 and 13.3 seconds resulted in some cases in a significant decrease in the aliased tone amplitude. The loss fraction is given by sinc(ϕ), where ϕ is the phase rotation in turns within the integration period, and sinc(x) = $\sin \pi \mathrm{x}/\pi \mathrm{x}$. This factor was calculated for each observation, using the observed phase wind, and the necessary correction applied to the observed reflected OH line amplitude.

The results of the measurements are shown in Figure 22.

None of the observations taken at the '1-dB' compression levels showed any detectable reflected OH emission, to a level of about 1%. At higher levels of compression, the strength of the reflected emission is generally a factor of about two greater than the model prediction.

10.2 Reduction of the Visibility

A drop in correlation coefficient of the source emission is expected, as reflected, incoherent emission is folded into the band of interest due to the 3$^{\rm rd}$order intermodulation mechanism. To determine the antenna-based loss, we used a model of the emitting region and the AIPS calibration program to estimate the antenna-based gains for both the 'tone-off' and 'tone-on' observations, for each of the three compression states. All the modified VLA antennas showed a drop in correlated power, while the unmodified ones remained constant, as expected. The results are shown in Figure 23.

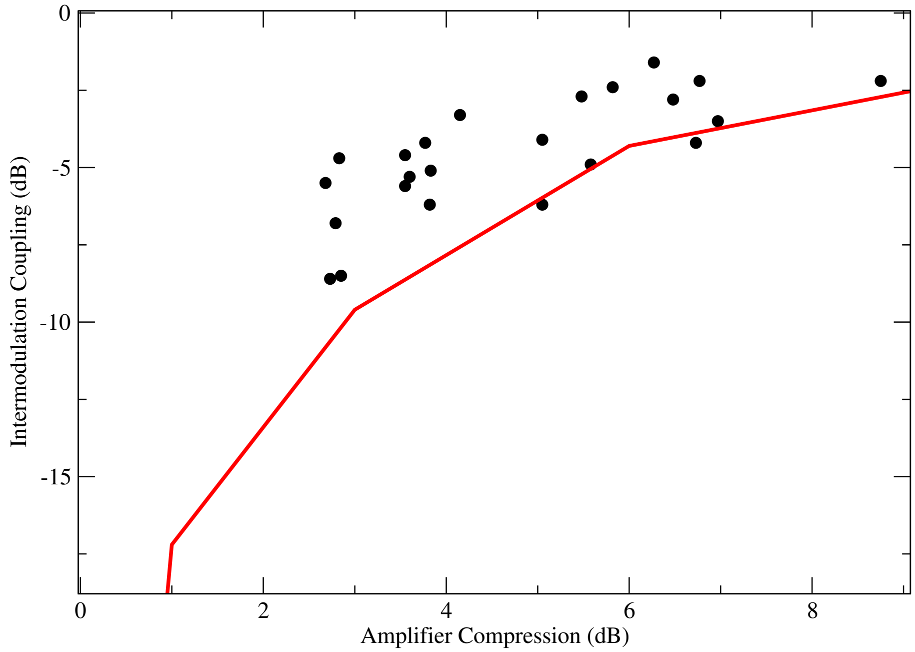

Figure 22: The observed coupling of the 3$^{\mathrm{rd}}$-order intermodulation of the OH line (black dots) and the model prediction (red). The actual coupling is stronger than the model predictions by a factor of about two for high compression levels. At the lowest, 1 dB, compression level, no reflected emission was seen, to a level of about 1% of the natural emission level.

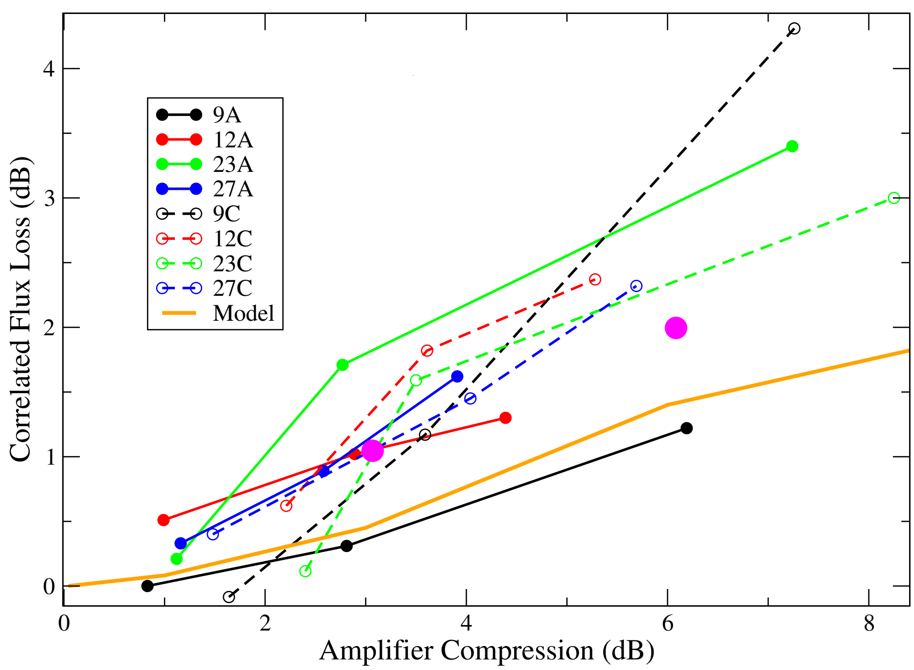

Figure 23: Showing the loss in correlated flux density as a function of the compression of the amplifiers. The yellow line shows the model prediction. The large magenta dots are the predicted losses when the observed coupling values shown in Figure 22 are used instead of the model predictions.

The visibility loss follows the expected trend, although there is considerable scatter between different antennas, and the loss in general is greater than the model predictions (in yellow). The under-prediction is probably because the reflected emission is higher than expected. The two large magenta dots are the expected loss if we take the observed coupling values shown in Figure 22, and apply the simple model described by equation 15.

The expectation that no significant change in system temperature due to the presence of the saturating tone was confirmed. Only one antenna-IF – 12A – showed any change in system temperature due to the introduced tone. For this antenna IF, the system temperature decreased point the intermodulation products become especially notable. Based on surveys of RFI power done so far, we expect the fraction of time that impulsive RFI will drive the receivers beyond the 1 dB compression point to be negligibly small. by 5K with the introduction of the 1 dB tone, was unchanged with the 3 dB tone, and increased by 11 degrees with the 6 dB tone. We have no explanation for this anomalous behavior.

10.3 Phase Rotation

We showed in Section 4 that the visibility phase of the 'reflected' OH emission will show a rotation, relative to the fringe-stopped phase tracking center, equal to

$$\dot{\phi}_3 = 2\omega_R \dot{\tau}_g + 2\Delta\dot{\phi}_R \tag{17}$$

where $\dot{\tau}_G = -\omega_e B_u \cos\delta/c$ is the rate of change of geometric delay due to earth rotation, $\Delta\dot{\phi}_R$ is the differential phase rate of the saturating tone, and ω_R is the angular frequency of the saturating tone. The first term accounts for the 'natural' phase rate appropriate for a source at the RFI frequency, while the second term accounts for any phase differentials due to the tone itself. For stationary interference, this latter term would be zero, and we would expect the reflected tone to show a phase rotation double the natural fringe rate at the frequency of observation.

However, in this experiment, the L6 tone that was used to put the receivers into compression includes a phase rotation normally used to fringe-stop the astronomical source in the BD channels. The observational setup chosen for the tone-on observations set the BD frequency to a value appropriate for an observation of 922 MHz (the L6 was set to 3090 MHz, and the Fluke synthesizer to 222.589 MHz). Because of the Hittite divide-by-two boards, the imposed phase rate was that appropriate for an astronomical source at a frequency of 461 MHz.

Because both terms are proportional to the natural fringe rate at 1422 MHz, the resulting phase rate of the 3^{rd}order intermodulation product can be written

$$\dot{\phi} = 2\left(\frac{\nu_R}{\nu_A}\right)\nu_f - 2\left(\frac{\nu_B}{\nu_A}\right)\nu_f \tag{18}$$

where ν_R is the RFI frequency, ν_A is the frequency of the observation, ν_B is the effective observing frequency of the BD IFs for which the saturating L6 was set, and $\nu_f = -\omega_e B \nu_A \cos\delta/c$ is the natural fringe frequency at the frequency of observation, ν_A. The sign of the second term is reversed because the effect of the 'false' fringe rate in the BD IFs is to partially compensate for the natural rate. For our experiment, $\nu_R = 1545$ MHz, $\nu_A = 1422$ MHz, and $\nu_B = 461$ MHz. Inserting these values, we predict an observed fringe rate of

$$\dot{\phi} = 1.52\nu_f = 0.0187u \quad \text{degrees/second} \tag{19}$$

where u is the E-W component of the baseline, measured in wavelengths at a frequency of 1422 MHz[5].

In Figure 24 we show the astronomical phase of the channel in which the reflected emission appears. When the tone is off, the amplitudes are constant, and the phase slowly winds, reflecting the rotation of the dominant source of emission about the phase tracking center. When the tone is on, the phase rotates rapidly, and the amplitude shows the beating between the true emission and the reflected tone. It is an elementary exercise to separate the two components. The result of this showed that the rate of change of phase of the tone emission matches the prediction exactly.

[5]We express this in this way because AIPS provides us this baseline component in these units.

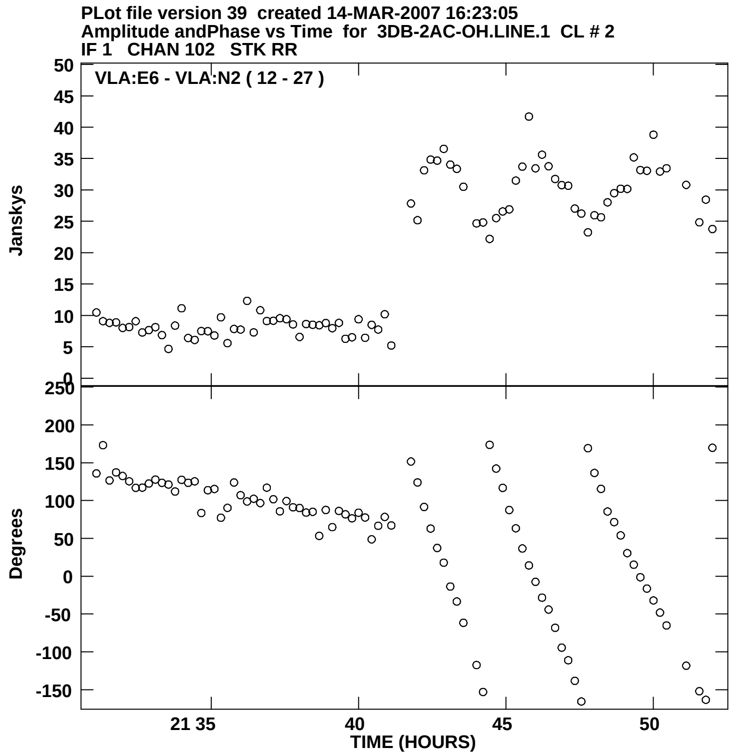

Figure 24: Showing the amplitude and phase as a function of time for baseline 12 – 27, when the compressing tone is off (left half) and on (right half). The sinusoidal variation seen in amplitude, and the notably non-linear variations in the phase are matched very well by a model of two complex vectors, one constant in amplitude and phase, the other constant in amplitude, but rotating at the rate given in Equation 19. The slow change in phase in the 'tone-off' state is due to the rotation of the source through the interferometer fringe pattern.

11 Folding over of the RFI Spectrum

The evaluation test on Antenna 23 described in Section 7 provided insights into the ambient RFI environment at the VLA. While this had little bearing on the experiment we were planning, they are recorded here for future reference.

Figure 25 shows the RFI spectrum across a broad frequency range seen by the L-Band receiver. The LCP and RCP channels are shown separately. The plots at the top show the frequency response in an uncompressed receiver. The spectrum analyzer used to take the data was in 'Max Hold' mode (we waited a minute or two for each trace in order to catch a reasonably good sample of transient events) with a 1 MHz resolution bandwidth. Note that although the cut-off frequency of the circular waveguide sections at the input of the L-Band receiver is about 1100 MHz, noise power in the bandpass as low as 600 MHz and extending as high as 2200 MHz is still seen. Much of the spectral power seen outside the normal L-band range must be generated in the LNAs and post-amps themselves. Thus, any RFI seen outside the optimum 1200-1800 MHz band must be very strong indeed.

The middle plots in Figure 25 shows the frequency response of the L-Band receiver when it is moderately compressed (blue traces) from our 1545 MHz tone, while the bottom plots show the spectrum when the system is in a heavily compressed condition (red traces). The light black traces show the non-compressed responses for ease of comparison. Based on the noise floor compression curve, the moderate compression case probably has the system near its 1 dB compression point. The overall frequency response has changed significantly and a number of 3^{rd}order intermod lines have appeared.

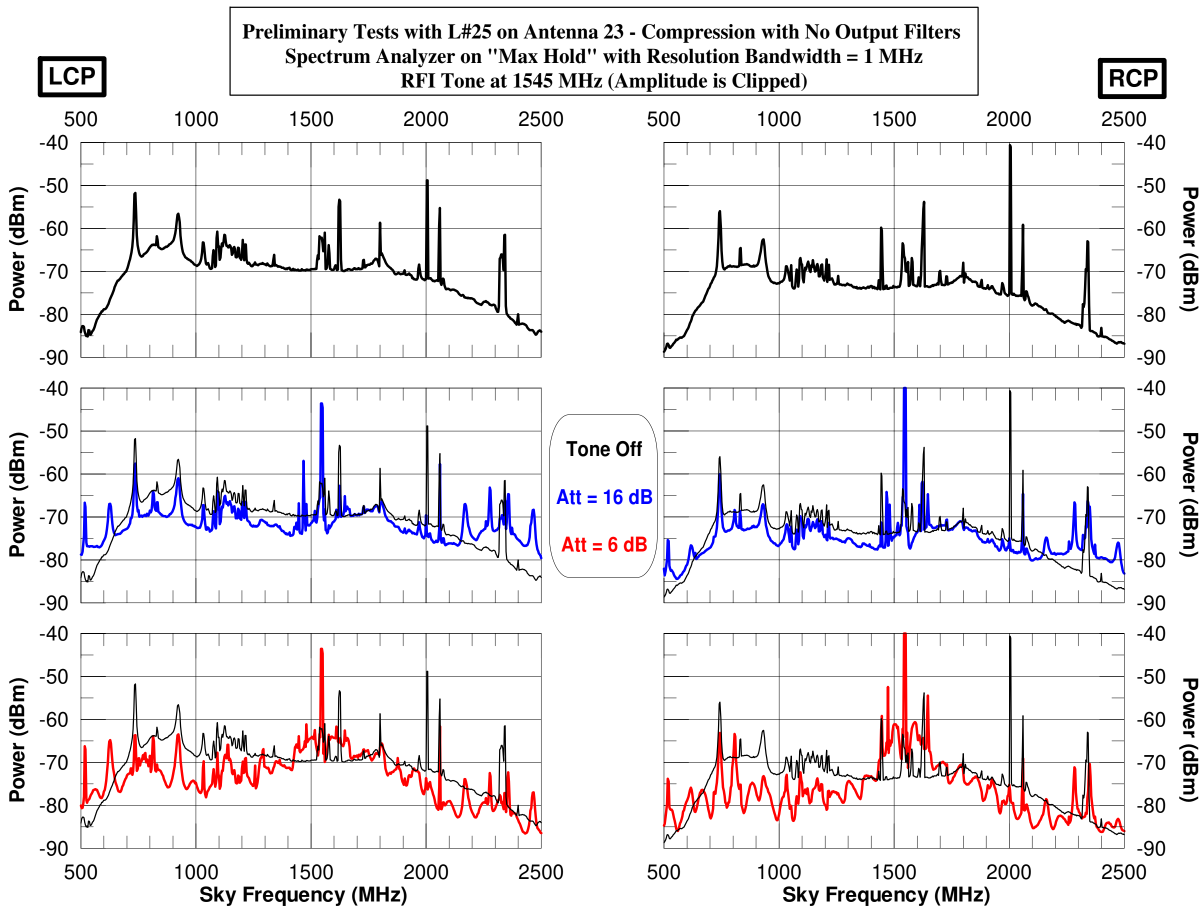

Figure 25: RFI Environment seen by L#25 on antenna 23, and the effect from a moderate to strong saturating tone at 1545 MHz.

In the strong saturation case, the noise floor in the 1400-1700 MHz is elevated (in a manner reminiscent of what we saw earlier in Figures 11and 12) and, more peculiarly, a ripple with a 60-70 MHz period now dominates the passband structure.

Figure 26 shows how strong RFI can cause other interfering signals to be folded across the output response of the receiver and appear in places of the spectrum that might cause the astronomer problems. This plot shows the output of the receiver RCP channel using the 1426 & 1675 MHz filter bank described earlier. The uncompressed condition is shown in blue while the reasonably strong compressed measurement is shown in red. There appears to be little compression of noise floor in the 1666 MHz band while the 1421 drops significantly (roughly 4 dB).

The birdy seen at 1445 MHz is of interest. When our RFI tone at 1545 MHz enabled, a response at 1645 MHz appears. This is the same type of 3rdorder product we are generating in the on-sky tests (i.e., 2x1545 - 1445 = 1645 MHz). Notice that there is also a birdy at 1465 MHz that appears when the system is compressed. This is from the RFI seen at 1625 MHz, which looks weak because it is sitting in the lower wing of the 1650-1700 MHz filter. However, inside the amplifier chain, it is causing a 3rdorder product (i.e., $2 \times 1545 - 1625 = 1465$ MHz). This real-life experience should cause us concern. As the RFI environment gets worse, and with the much wide receiver bandwidths of the EVLA, these sorts of false ‘reflections’ may be commonplace.

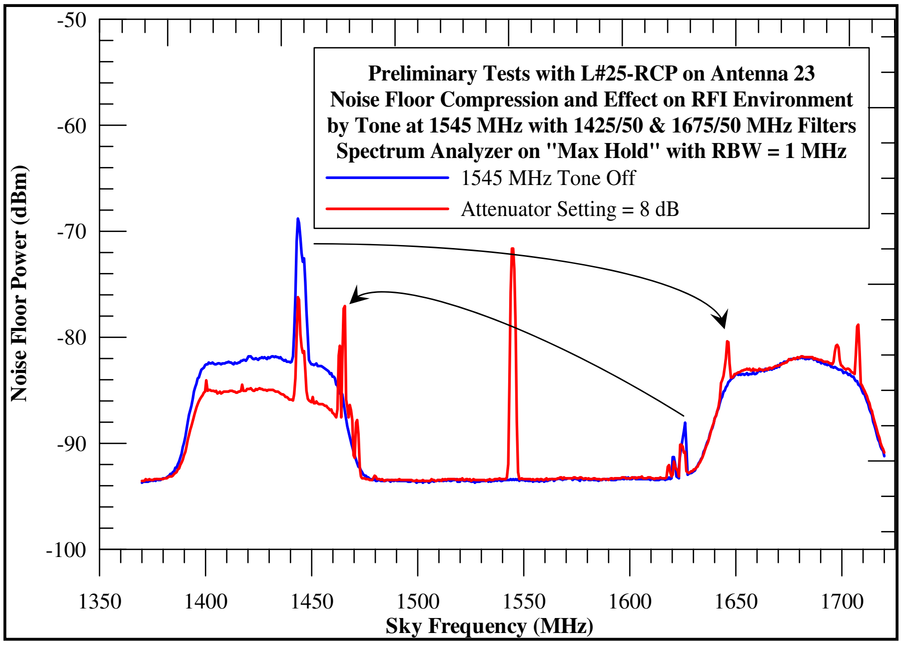

Figure 26: Examples of aliased ambient RFI lines seen on Antenna 23 from a strong 1545 MHz tone. The response in red (tone-on state) at 1645 MHz is due to strong RFI at 1445 MHz reflected by the 1545 MHz tone via the third-order intermodulation. Similarly, the responses near 1465 and 1470 MHz arise from emission at 1625 and 1620 MHz. The responses seen near 1698 and 1708 arise from emission outside the lower band filter response.

12 The EVLA's L-Band Amplifiers

At this point a few words about the design of the EVLA L-Band receiver are appropriate. This next-generation receiver will be outfitted with new broadband balanced-amplifier gain blocks developed by NRAO's Central Development Lab. Amplifiers with excellent input impedance matching are necessary in order to meet the polarization purity requirements imposed by the use of a 90-degree hybrid coupler following the orthomode transducer to derive circular polarization. Since wideband cryogenic isolators are unavailable at L-Band, broadband balanced amplifiers offer the only viable solution. At the request of the EVLA Project, the CDL developed both a new 1-2 GHz low-noise balanced-amplifier as the first gain block and a high-power balanced-amplifier for the second gain block. Two cascaded gain blocks will be used per channel (for a total of 4 amplifiers in each dewar). The first gain block uses custom InP transistors to achieve a noise temperature of about 4K and an input return loss of better than 15 dB. The second gain block uses commercial HFET's to achieve a $P_{1dB} > +13$ dBm but with a noise temperature closer to 20K. Both blocks have $\sim$ 18 dB of gain.

This low-noise, high-power amplifier configuration provides an ideal compromise between sensitivity and headroom. It ensures that the cryogenic amps won't saturate before the warm post-amps from the effects of the RFI currently seen at the VLA site. The scheme allows for the possibility of adding a switched filter bank between the dewar output and the room temperature post-amps (which will be the first to saturate) if necessary in the future. It also allows for the option of adding cooled rejection filters between the cryogenic gain blocks should the RFI environment worsen dramatically. Should extremely strong interference appear, the ideal solution is to place a high TC super-conducting filter (which, of course, does not yet exist) as early in the signal path as possible. If it is located before the first gain stage, the resulting increase in receiver temperature would be unacceptable. Placing a filter between the gain blocks ensures the best tradeoff between minimizing the degradation in sensitivity and increasing the dynamic range.

Eight EVLA interim L-Band receivers have been built as of April 2007 and are placed on the EVLA antennas. These are existing VLA systems that have been modified with the octave bandwidth gain blocks along with a hybrid coupler circular polarizer. Seven if these retain the old VLA narrowband 1.2-1.8 GHz orthomode transducer, but nevertheless allow us to characterize the performance expected from the final EVLA version. In December 2004, compression tests were carried out on a VLA L-Band receiver (L#11), which utilized old-style cooled single-ended amplifiers, and the first receiver (L#32) to be outfitted with the balanced-amplifier blocks. The 1 dB compression point was measured for each receiver using a swept CW signal. Without the warm post-amps, while L#11 had an input saturation power level of -34 dBm while L#32 had -22 dBm – 12 dB improvement in headroom. With the post-amps included, L#11 saturated at -48 dBm and L#32 at -47 dBm, which indicates that the post-amps are the limiting factor. It turns out that these 1.2-1.8 GHz amps are surprisingly weak, with a P1dB point of only about -5 dBm). The new EVLA receiver design will utilize higher power post-amps with a $P_{1dB} > +15$ dBm.

In order to avoid the adverse affects of non-linearities which may arise in amplifiers, mixers, or diode attenuators that are driven into compression, the design of all the EVLA front-ends, and the subsequent IF modules in the signal path, attempts to ensure that the standard operating point (i.e., the power level when observing 'cold' sky) is well below the saturation point of any active device. This is a delicate balance since the signal must also be well above the noise floor of the various amplifiers in the RF/IF chain to ensure the overall system temperature is not degraded. The headroom specification in the EVLA Project Book calls for all the active components to be operating at least 20 dB below their 1% compression point (which corresponds to being about 32 dB below their 1 dB compression point). The major impact of insufficient headroom will only become a major limitation in the presence of strong interference, where unwanted harmonics and intermodulation products generated in the amplifiers and mixers may cause spurious signals to arise. The 20 dB headroom spec should provide adequate dynamic range to mitigate the effects of both existing and future RFI. By utilizing high-power cryogenic balanced-amplifiers and post-amps, the new EVLA L-Band receiver system will achieve a much improved headroom specification over the existing VLA front-end. When looking at cold sky, the EVLA receiver should have over 40 dB of headroom. Interference like the DME signals from airplanes, which are expected to increase the instantaneous integrated power seen in the 1-2 GHz band by about 20 dB, should be handled comfortably by the receiver. Evidence that this is so comes from EVLA antenna 14, which has the prototype OMT installed, thus exposing the amplifiers to the full range of the aernautical band's signals. Comparison of this antenna's sensitivity and performance in the standard observing bands with that existing prior to the installation of the new OMT has shown no notable change.

13 Summary

We have conducted a carefully calibrated experiment in which four VLA L-band receivers were put into known degrees of amplifier compression by the introduction of a noise tone at 1545 MHz. The resulting compression reflects, via the third-order intermodulation mechanism, the strong OH maser emission at 1665 MHz to a frequency of 1425 MHz, adjacent to the HI absorption line.

We have measured the degree of coupling of this reflection mechanism through astronomical observations of W3OH, an object which contains both a strong OH maser, and a nearby continuum source, in front of which lies neutral HI. The observations show the coupling efficiency of the intermodulation rises sharply with degree of compression, such that at 1 dB compression, the reflected emission is seen at $\sim$1% of its true intensity, while by 3 dB compression, it has reached $\sim$25% of actual strength.

The third-order intermodulation mechanism introduces a phase coupling into the reflected signal which, for a non-moving external saturating signal, will cause the interferometer phase of the reflected signal to rotate, relative to the astronomical source of interest, at twice the natural fringe rate. This relatively fast rotation is very useful in decorrelating the reflected emission, and preventing the appearance of a 'ghost' image of the reflected emission into the desired image.

The reflected signal, however, will introduce extra noise, resulting in a reduction of the correlation coefficient of the desired emission. This reduction cannot be recovered by normal system temperature calibration methods, as the reflected signal has the synchronous calibration noise embedded. The loss could be corrected through normal off-source calibration, provided the saturation is unchanged between calibrator and target.

The saturating signal converts power out of the fundamental frequency into higher harmonics, resulting in an effective loss of gain for both the tone power and the astronomical power. Curiously, the loss for (assumed weak) astronomical power is as high as twice that of the tone power. This gain loss in noise power would be automatically corrected for by standard gain stabilization techniques, provided the saturating signal is stable over the integration time of the gain monitors. This situation, however, is unlikely to pertain for most observations, where the RFI is sporadic and/or highly time variable. In this case, the reduction of gain will be uncompensated for, and a potentially large loss in apparent interferometer visibility will result.

The best defense against these RFI-induced effects is to design radio astronomical amplifier systems to provide the largest possible range between the cold-sky operating power and the 1-dB compression point, which is where the amplitude scale errors become truly noticeable. It is noted that higher head-room must inevitably be traded off against better sensitivity, so that knowledge of the expected RFI power spectrum is essential for optimum design of the amplifier chain. The goal must be to design the system so that the head-room is set such that a negligible fraction of the observing time is spent in a state of compression exceed ~1 dB.

The EVLA L-band design has incorporated these goals and constraints. Initial tests with the first antenna to be outfitted with the new wide-band OMT shows no noticeable loss in sensitivity.

14 Acknowledgements

This experiment was facilitated by the expert assistance of Brent Willoughby and John McClenden. We thank Tanner Oates for his software for capture spectrum analyzer data. We are especially grateful for the assistance provided by Dan Mertely in outfitting, calibrating and adjusting the four modified antennas, and for the discussions held with Ken Sowinski, which were instrumental in understanding some of the interesting phenomena generated by this experiment.

References

[1] Rick Perley. EVLA Memorandum #66, The Effect of Electronics Non-Linearity on Closure and Imaging. August 2003

[2] Rick Perley and Bob Hayward. EVLA Memorandum #79, The Effect of Strong CW Signals on Interferometer Calibration, Closure, and Imaging. July 2004

[3] A. Harris. ALMA Memorandum #307, Precision Radiometry and the APHID 22 GHz Water line Monitor. May 2000