Why wide bandwidths are good for geodetic VLBI

1) Maximize signal-to-noise ratio SNR =

$$\frac{(0.5 - 1) \times (\text{correlated flux density}) \times (2 \times \text{total bandwidth} \times \text{time})^{1/2}}{(\text{SEFD}_1 \times \text{SEFD}_2)^{1/2}}$$

where SEFD = (time-averaged system temperature) x (Jy/K)

2) Minimize error in multiband delay = 1 / (2π x SNR x RMSfreq)

where RMSfreq^2 = average value of $(\text{RF freq} - \text{mean RF freq})^2$

Narrowband cross-correlation –
Delay measurement is ambiguous:
Which cycle at station A corresponds to which cycle at station B?

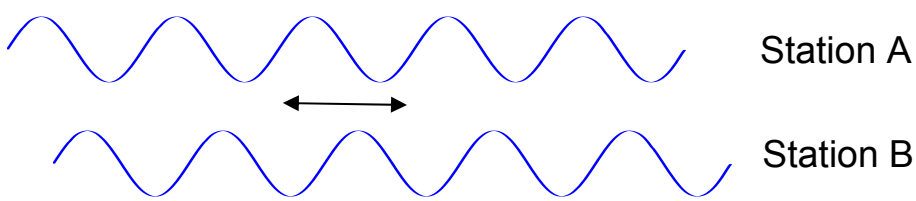

Wideband cross-correlation –
Unambiguous delay measurement

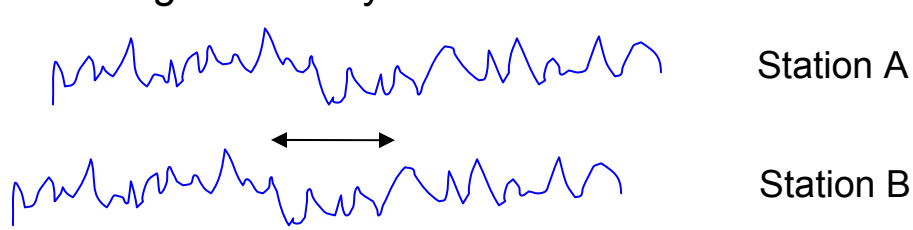

Effect of RFI on geodetic VLBI

The primary effect of RFI is to increase the effective system temperature in one or more frequency channels. This additional noise, which is usually uncorrelated between stations, has the following effects:

- decreases the SNR, and thereby degrades the delay precision
- can lead to systematic shifts in the estimated group delay
- can cause bandwidth synthesis to fail if enough frequency channels are severely affected
- adversely affects source mapping when RFI is variable.

Except in rare cases or in short-baseline geodetic sessions (e.g., “ties” experiments), correlated RFI signals between two stations are not found at the correlator. There are several reasons for this lack of correlated RFI:

- Except for satellites, most RFI sources cannot be seen at more than one terrestrial VLBI station, due to the wide geographical separation between stations.
- If one antenna is pointed at or near a satellite, the other antennas in the network will generally be pointed well away from it.
- An RFI source visible at two stations generally has a different delay and delay rate from the celestial source they are tracking, and so any correlated RFI signal will be strongly rejected at the correlator.

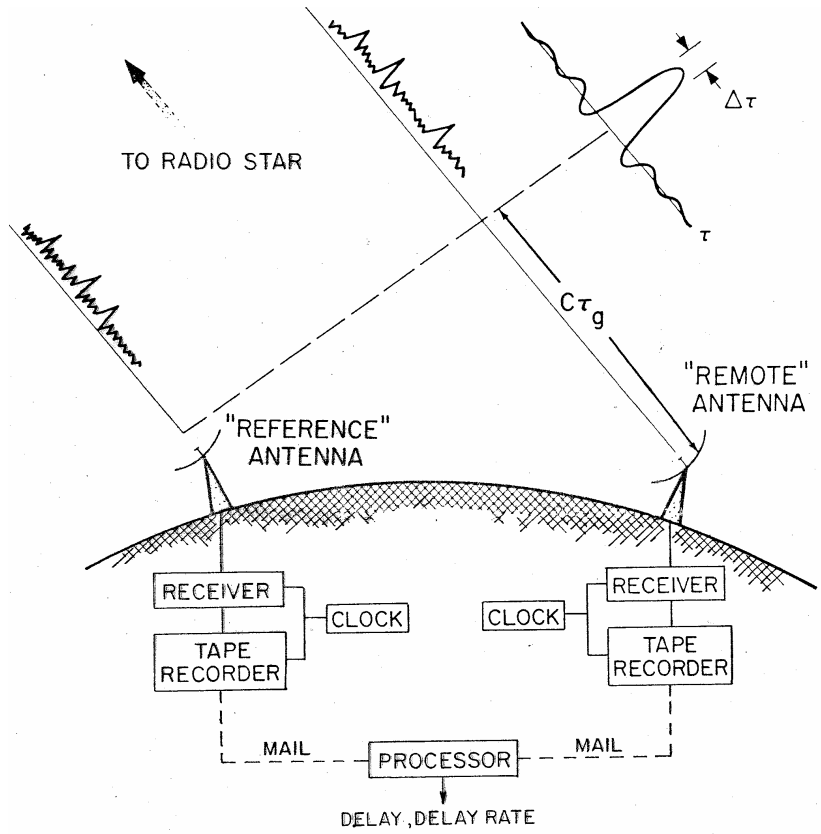

Effect of RFI on group delay observable

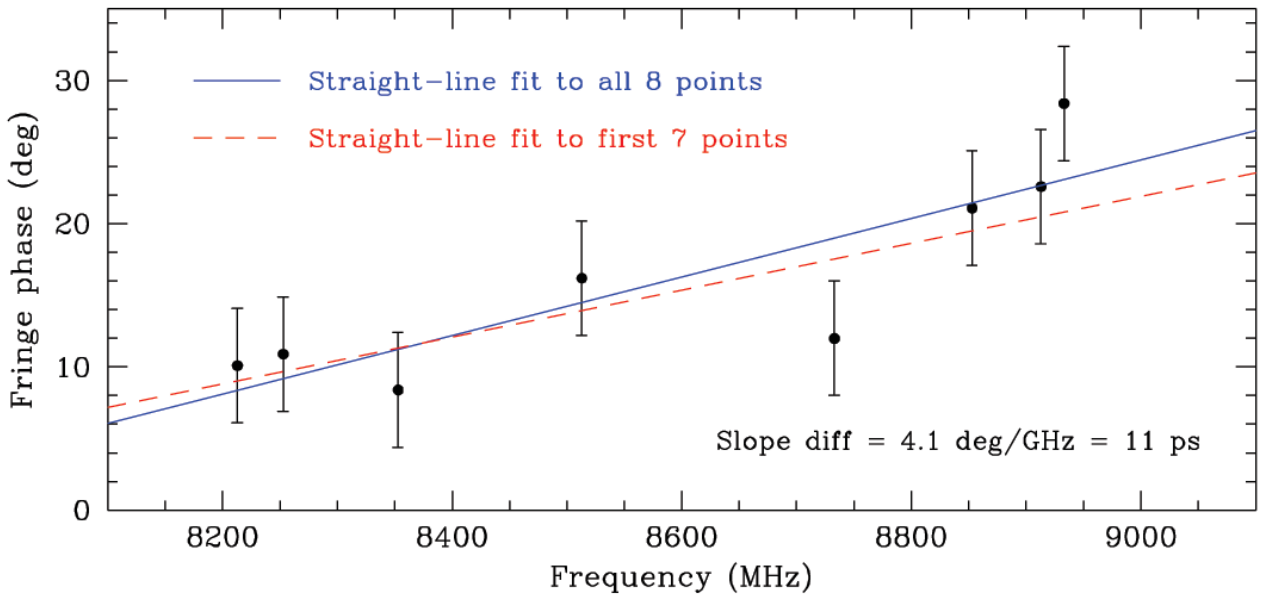

Group delay is estimated from the frequency dependence of the fringe phases. Its value is the slope of the straight line that best fits the phase plotted vs. frequency.

As the system temperature due to RFI (or any other cause!) rises in one frequency channel, the “weight” of that channel relative to the other channels will decrease in the fitting process.

If the channel with RFI has a fringe phase that is systematically biased (i.e., stands out) relative to the other channel phases, then as the system temperature in that channel varies, the group delay estimate will vary correspondingly.

The plot above shows how the group delay estimate can change between the two extreme cases (a) no RFI in channel 8, so the phase is included in the fit with full weight, and (b) infinitely strong RFI, so the phase is completely “down-weighted”, and is effectively excluded from the fit.

Simulations similar to the above demonstrate that the delay precision can be significantly degraded, by >1 ps, when RFI raises Tsys by >10%. See paper by Dave Shaffer in IVS 2000 General Meeting proceedings.

How strong must RFI be to cause a problem?

- If RFI power > 10% of system power in any baseband channel…
  - SNR is reduced
  - Group delays will be less precise, and may be biased significantly

- If RFI is coherent with maser, and its frequency is an integer number of MHz in RF or IF…
  - Phase calibration will be degraded if spurious signal is stronger than –50 dB relative to phase cal

- Even if the frequency of the RFI is outside the RF or IF observing band, it may be strong enough to compress the gain of amplifiers in the receiver or DAT/DAR at all frequencies, in which case…
  - Same problems as above, when RFI in-band power > 10% of system power

- For incoherent RFI, what matters is the fractional increase in total power over the full frequency range of the baseband channel. If the RFI raises Tsys by 100% over a 20-kHz bandwidth, but the channel is 2 MHz wide, the fractional increase over the entire band is only 1%, which is not a problem….unless you're doing spectral line VLBI!

Sources of RFI

External to VLBI system

- Usually originates at RF frequencies (2.2-2.4 GHz and 8.2-8.9 GHz)
- Might be at image frequency, if RFI is strong enough to overcome image rejection in receiver or VC/BBC
- Can come in at IF frequencies (100-1000 MHz) if RFI is exceptionally strong, especially if an IF cable has a broken shield or bad connector
- Potential interferers:
    - Satellites, particularly at S-band (DARS, DSCS – Defense Satellite Communications System)
        - For satellite orbital elements, see http://www.celestrak.com/NORAD/elements
        - For orbit tracking software, freeware and cheapware are available at http://www.amsat.org
    - Local TV and radio stations, and TV links (S-band especially)
    - Radar

Internally generated RFI

- RF or IF amplifiers may oscillate
    - LNAs are especially prone to oscillation, which may occur at high (RF) or low (IF) frequency
- Formatter generates harmonics of 4, 4.5, and 5 MHz
- Set LOs of unused BBCs to frequencies outside observing band

Identifying RFI

Spectrum analyzer is most efficient instrument to find narrowband RFI.

- The narrower the resolution bandwidth, the more a narrowband signal sticks out above the noise.
- But: a spike 30 dB above the noise may or may not be a significant source of RFI, when it's averaged over the full baseband frequency range; it depends on what the resolution bandwidth of the analyzer is set to. Don't be alarmed by every little blip on the analyzer.
- If RFI is very narrowband but shifts in frequency between analyzer sweeps, it may be radar.
- A single-frequency carrier is probably due to amplifier oscillation or pick-up of a signal generated locally by the VLBI equipment or other control room equipment.
- Rapid amplitude variations over tens of kHz to many MHz point to AM, FM, or TV.
- Flat-topped, constant-amplitude signals >1 MHz wide are usually data transmission, often from satellites.

A VLBI backend can be used as a poor man's analyzer, by stepping VC/BBC frequencies across observing band and measuring power levels.

System temperatures are sensitive to RFI, but it's not always easy to spot in the data, as Tsys may be varying from scan to scan due to legitimate causes like changing tropospheric antenna temperature and ground pickup.

Fringe and phase cal data from the correlator can be used to identify RFI. But don't count on the correlator to find it for you!

What can be done about RFI?

- To reduce out-of-band interference, install filters with sharper cutoffs.
- Try to get cooperation from agency operating interferer, e.g., by adjusting transmission schedule to avoid geodetic sessions.
- Change observing frequencies to avoid persistent RFI.
    - Has been done to accommodate RFI at Matera, Medicina, and Westford, among others.
- Avoid observing in direction toward interferer.
    - Will have negative effect on geodetic results
- If RFI is internally generated….
    - Fix the oscillating LNA!
    - Fix the broken shield on the IF cable coming into the rack!

Digital Audio Radio Service (DARS)

- DARS is a satellite broadcast system designed to provide audio programming primarily for motorists.
- Transmission frequency band over North America is 2320–2345 MHz.
- Two private companies provide DARS in North America:
  - Sirius Satellite Radio has satellites in elliptical, inclined, geosynchronous orbits that are transmitting at 2320–2332.5 MHz.
  - XM Satellite Radio satellites are in geostationary orbits and transmit at 2332.5–2345 MHz.

DARS signal picked up in far sidelobes of Westford antenna

S-band IF spectrum is shown.

Add 2020 MHz to convert frequency to RF.

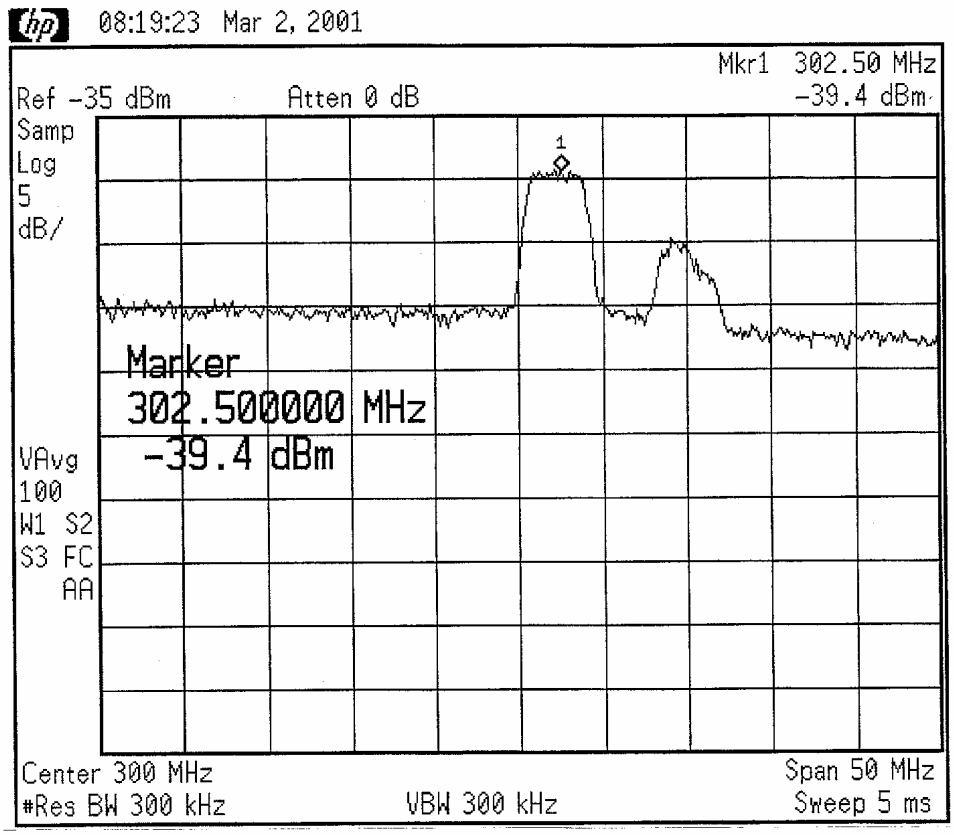

Example of RFI so severe that RF gain is reduced outside of RFI frequency range

Westford S-band IF spectrum on 26 August 2003

Antenna pointed at 179.6° azimuth, 2.5° elevation.

Add 2020 MHz to IF frequency to get RF frequency.

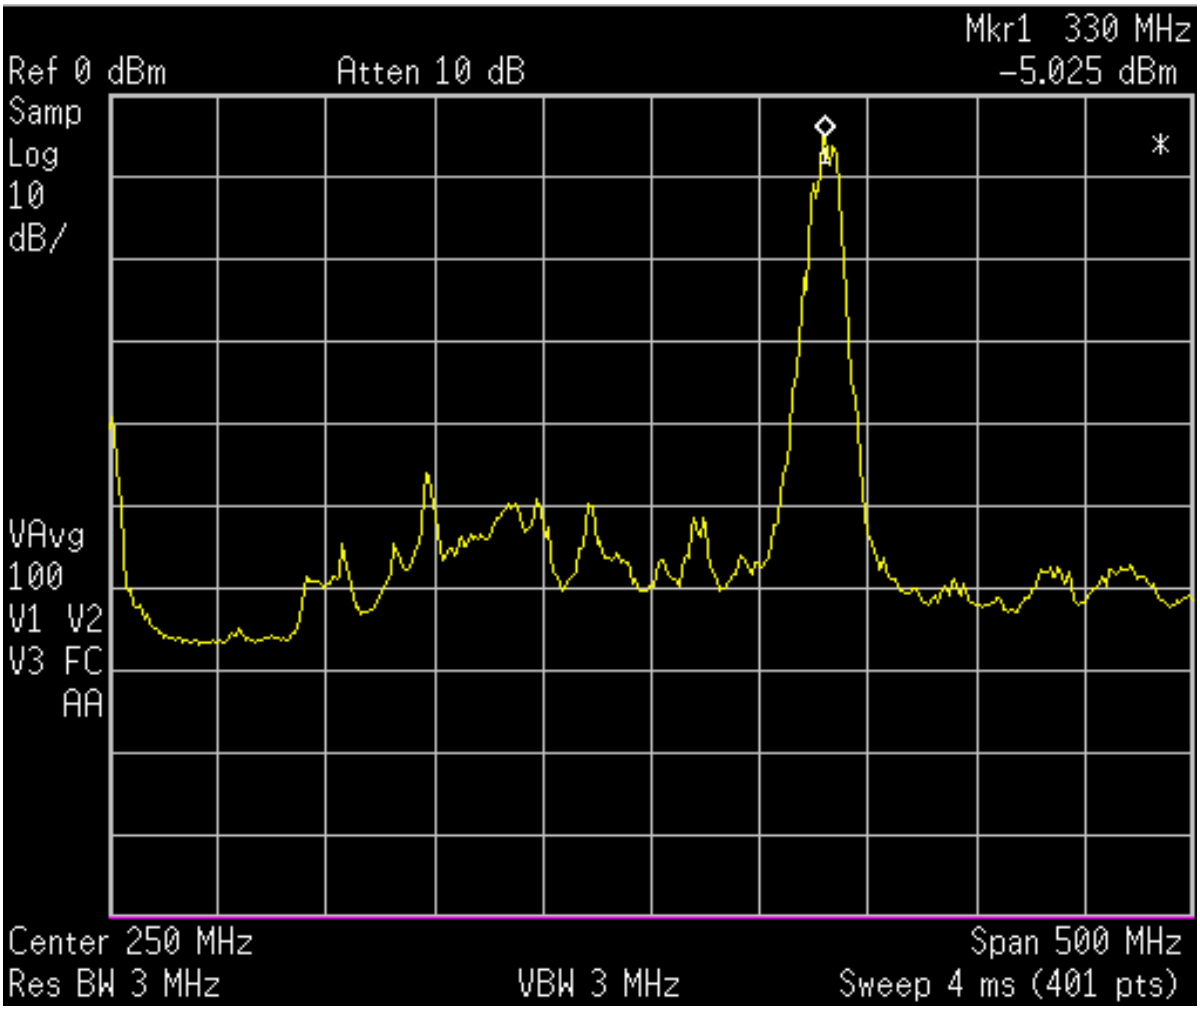

In absence of RFI, average power level at mid-band is approximately −50 dBm over 3 MHz RBW, i.e., halfway down the plot. At this antenna position, however, the power level is below this level at most frequencies.
→ System gain is being compressed due to the strong interfering signal.

Westford S-band IF spectrum
with higher frequency resolution over 320-345 MHz

Antenna pointed at 178° azimuth, 4.5° elevation.

Add 2020 MHz to IF frequency to get RF frequency.

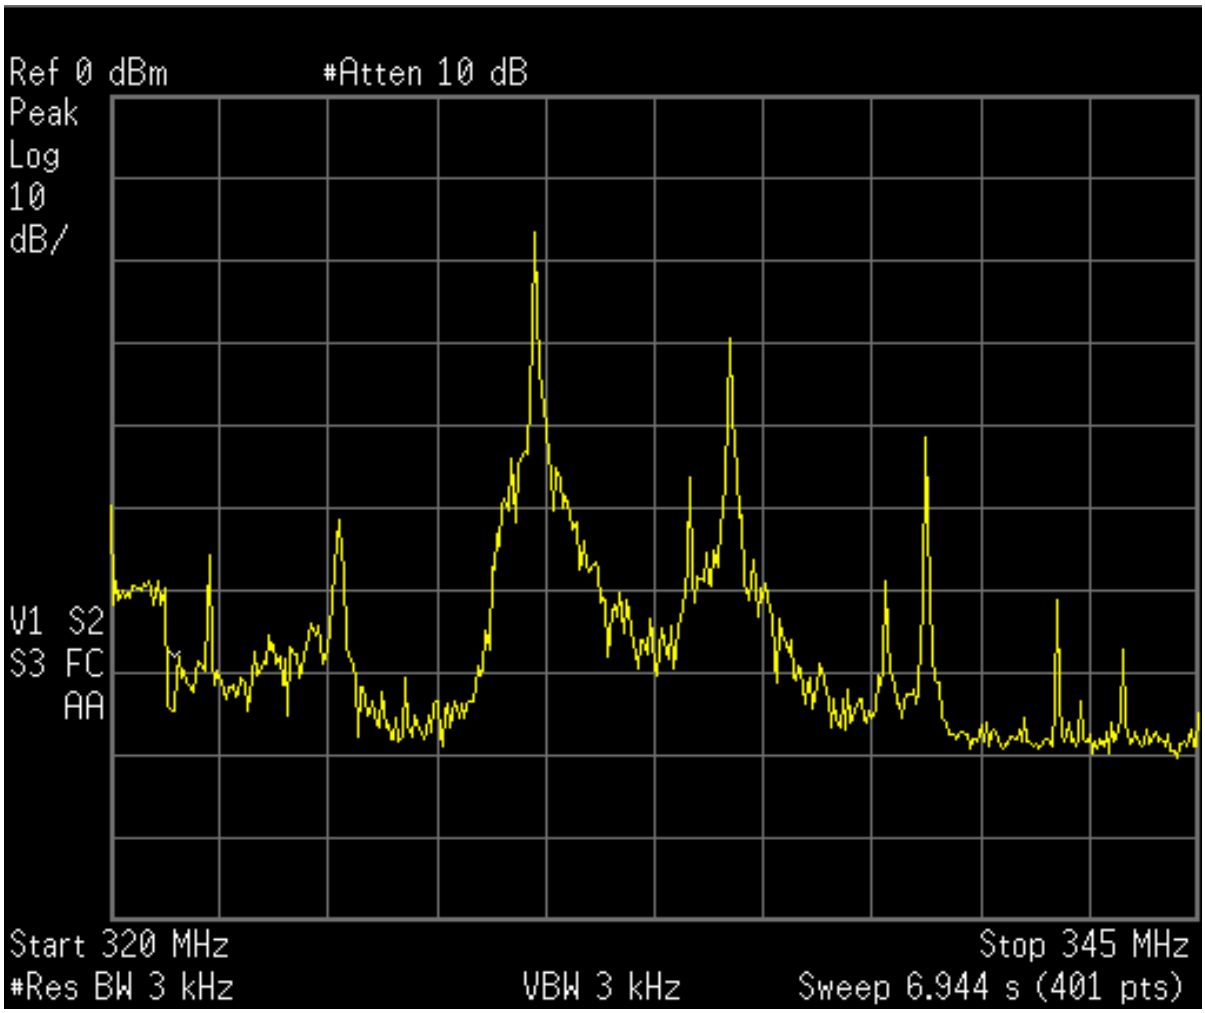

Example of severe RFI in several S-band channels:
Effect on measured system temperatures

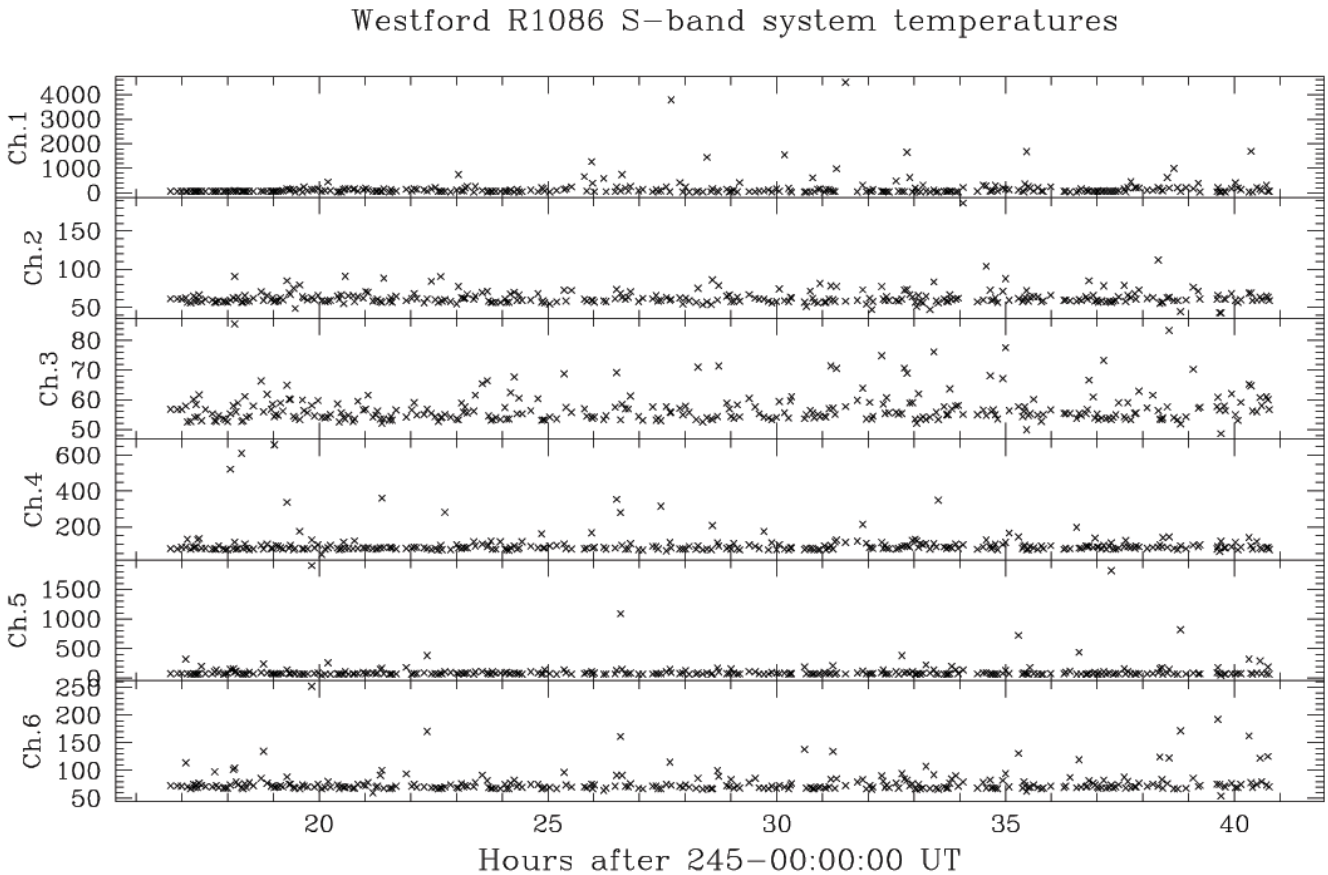

When RFI is strong and varies significantly on the time scales on which system temperature measurements are made (namely, the interval between the cal-on and cal-off total power measurements, which is typically of order seconds), system temperature measurements can be seriously compromised. In such cases, a better estimate of the effect of the RFI on the system sensitivity is provided by the raw total power counts, an example of which is provided on the next page.

Example of severe RFI in several S-band channels:
Effect on baseband total power levels

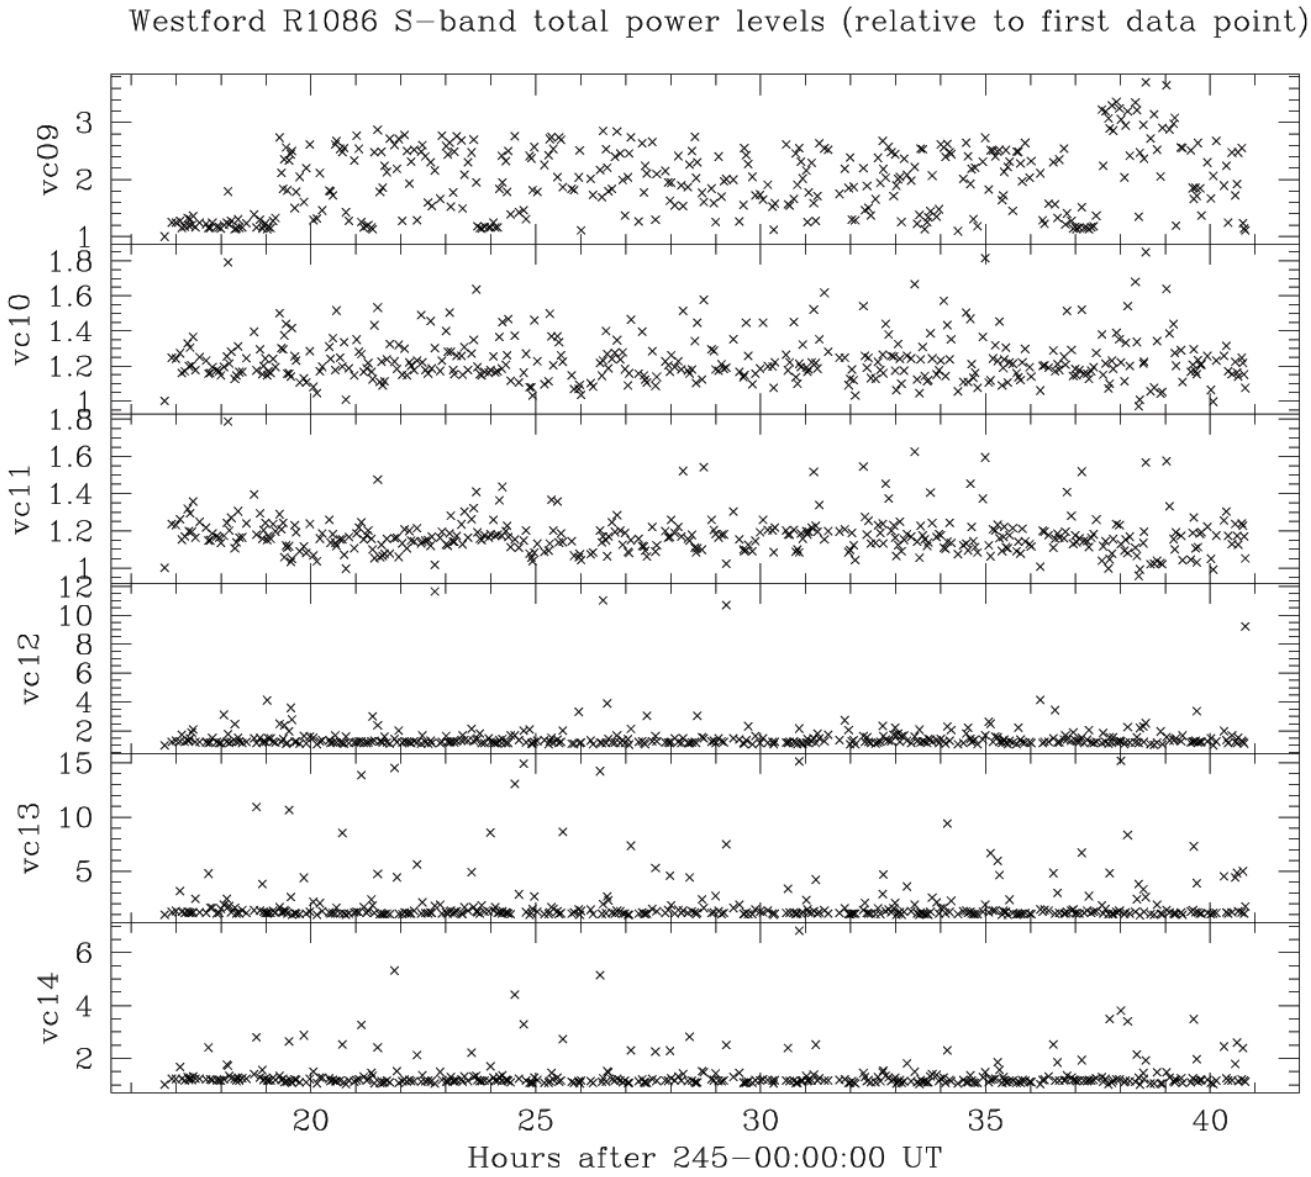

Fringe plot with RFI in one frequency channel

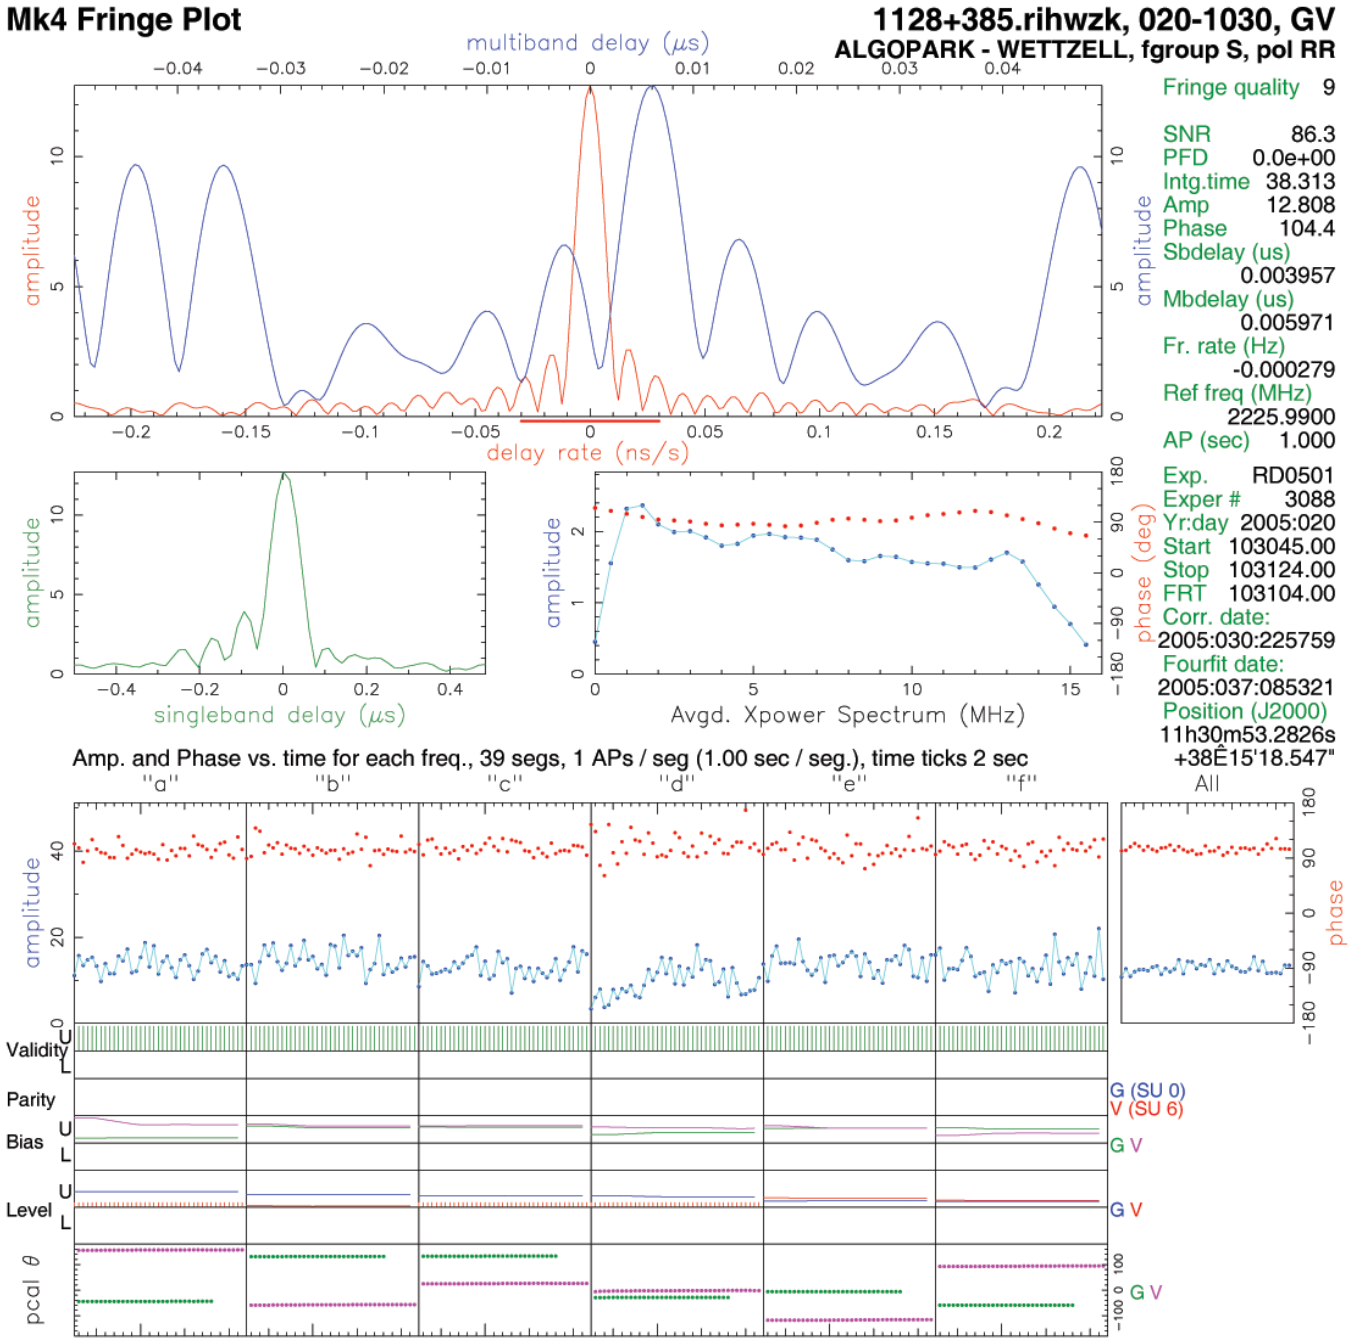

Can you spot the RFI?

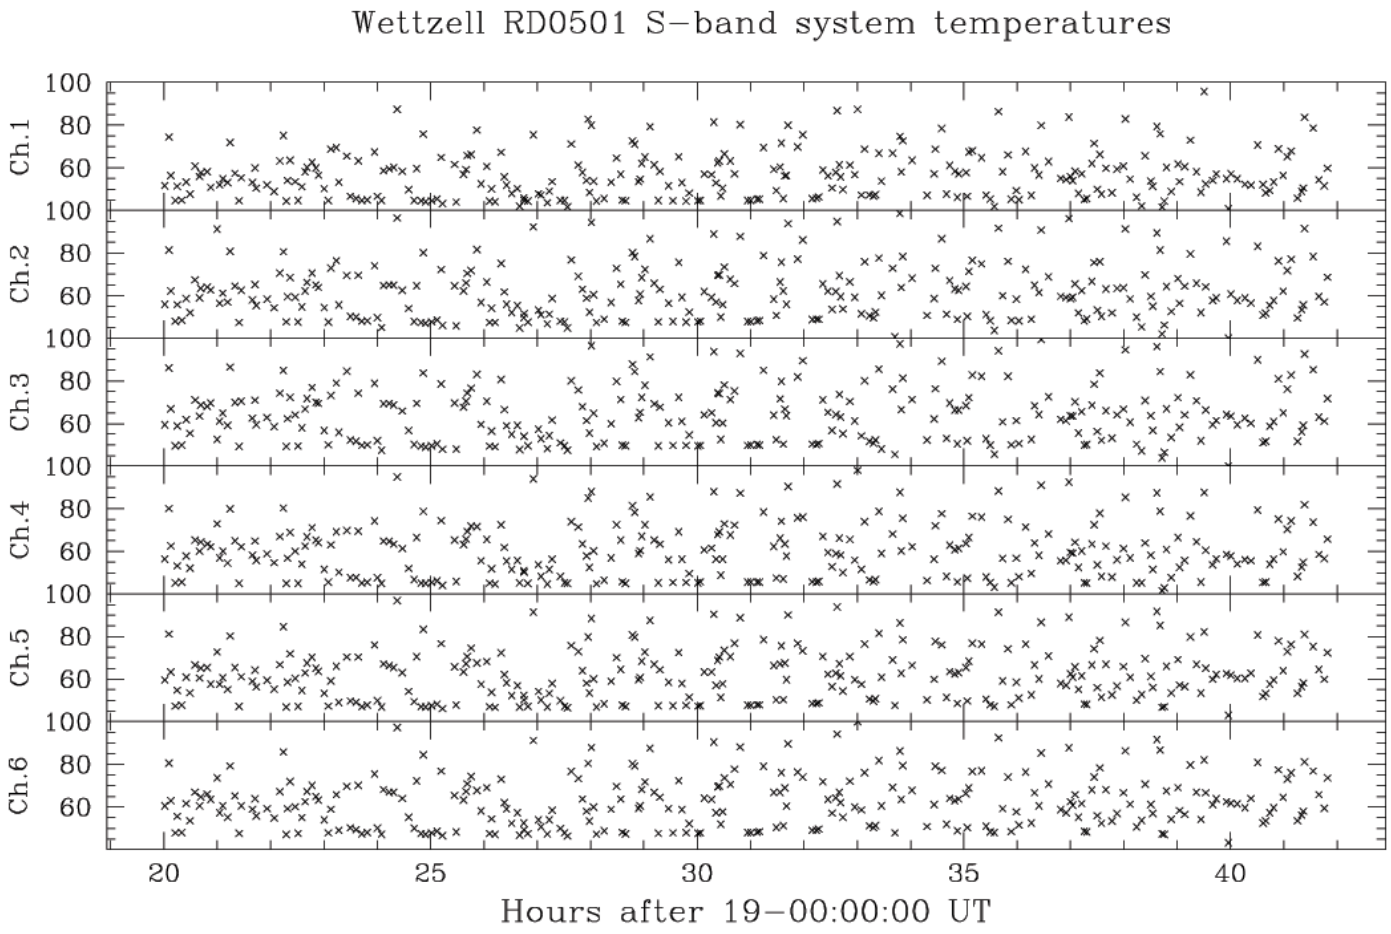

It’s a bit easier to see in the time series
of the phase cal amplitudes.

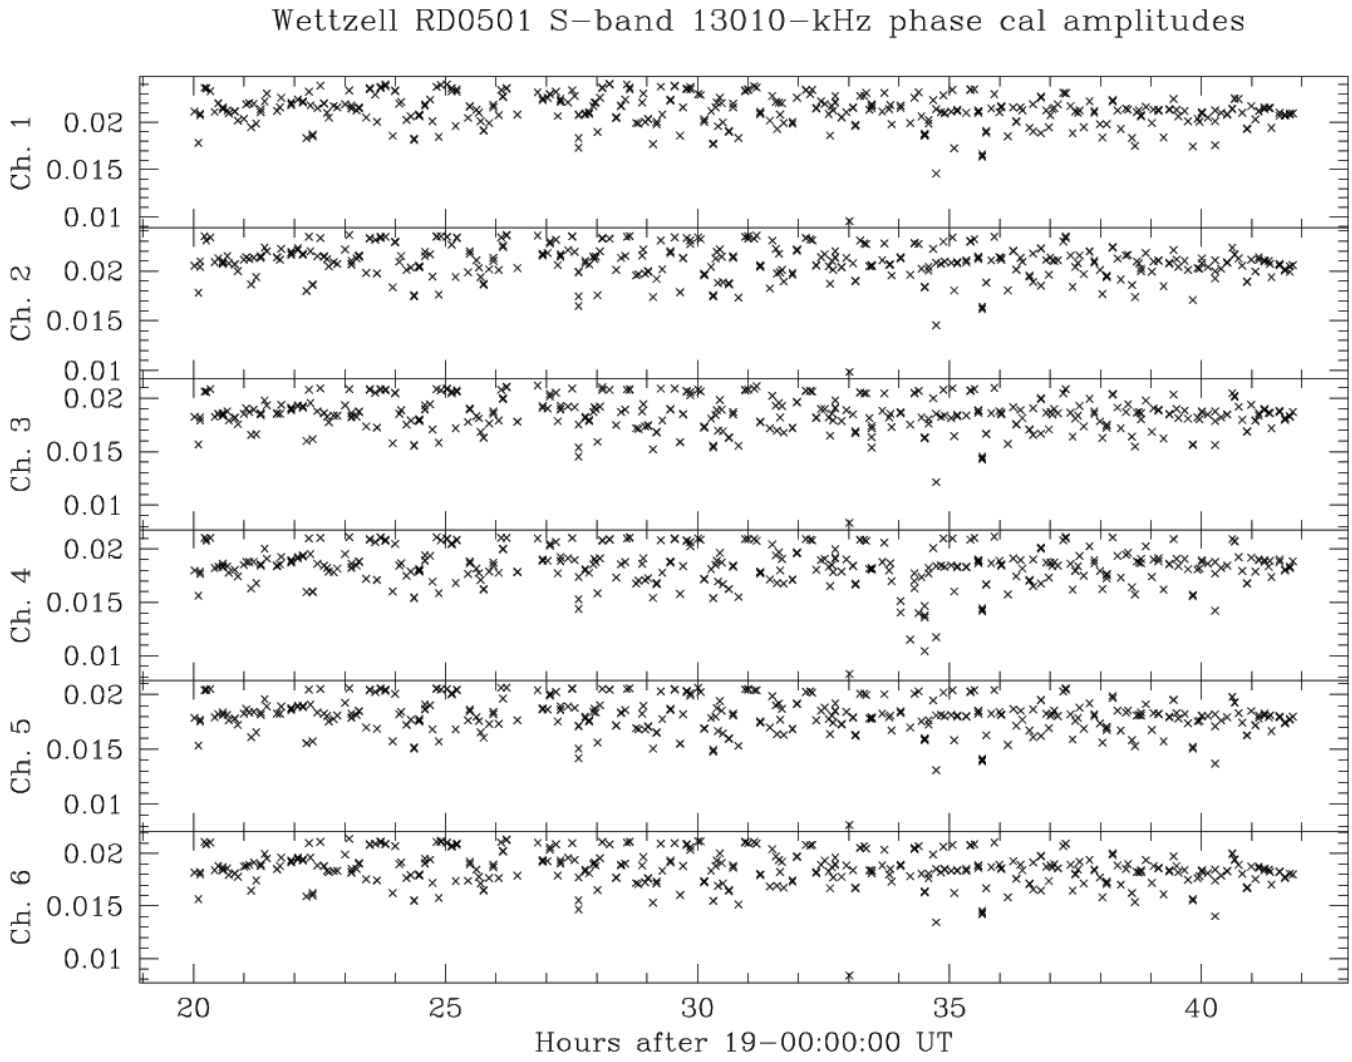

RFI is trivial to spot in the ratios
of the phase cal amplitudes in different channels.

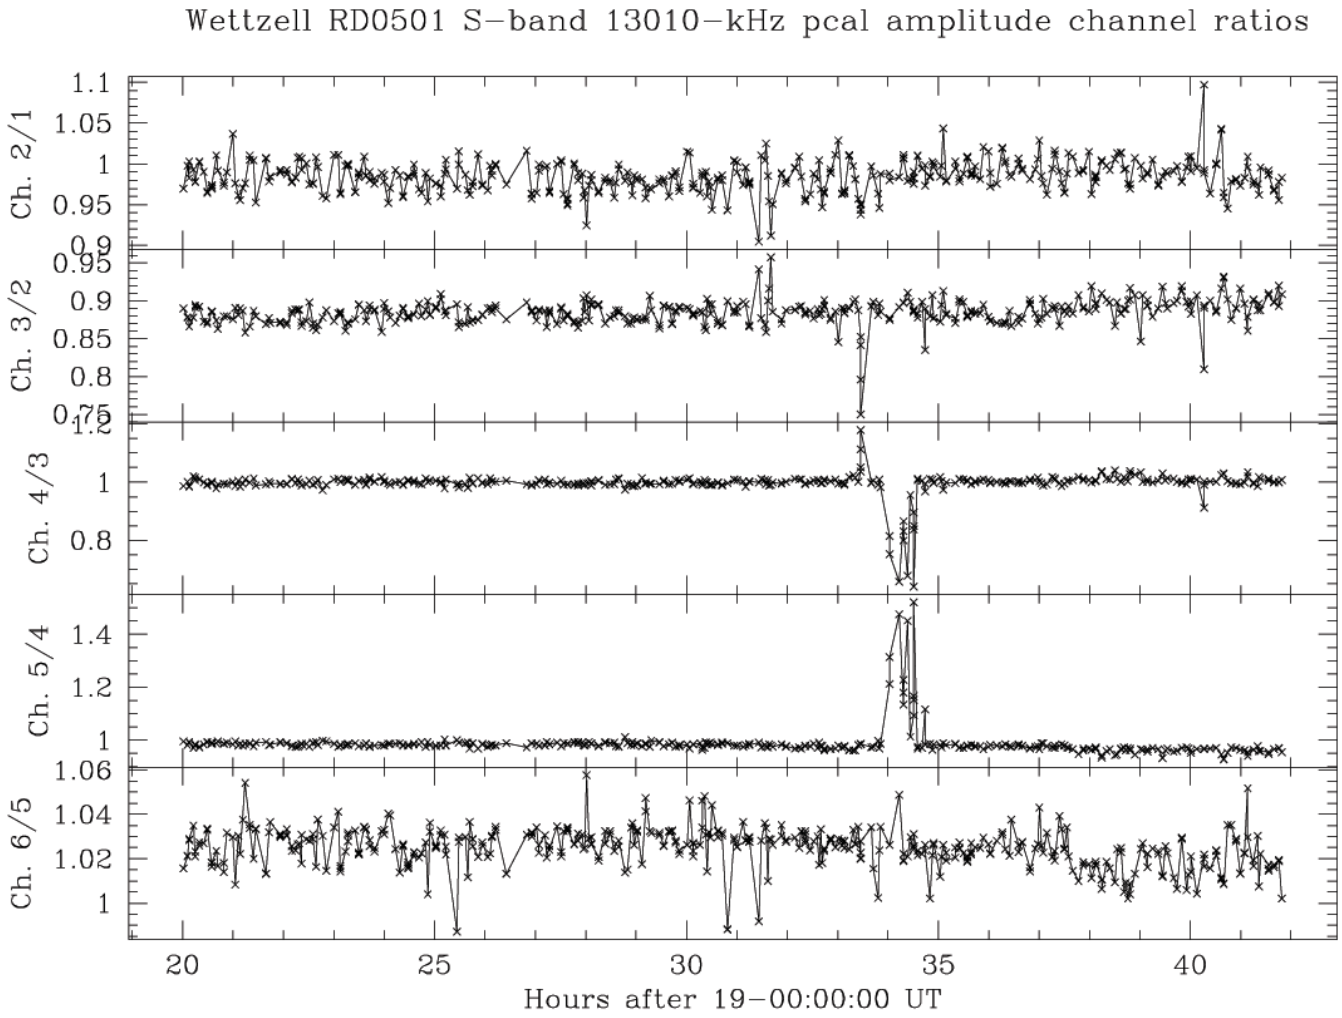

Time-variable RFI is more readily evident in ratios of parameters (e.g., phase cal amplitude or system temperature) from different channels, than from the parameter of a single channel, because common-mode variations due to weather or ground noise pickup, for instance, cancel out in the ratios.

RFI generally does not bias fringe or phase cal phases unless it is strong enough to saturate the electronics.

The reduction in SNR caused by RFI will, however, increase the fringe phase scatter in affected channels.

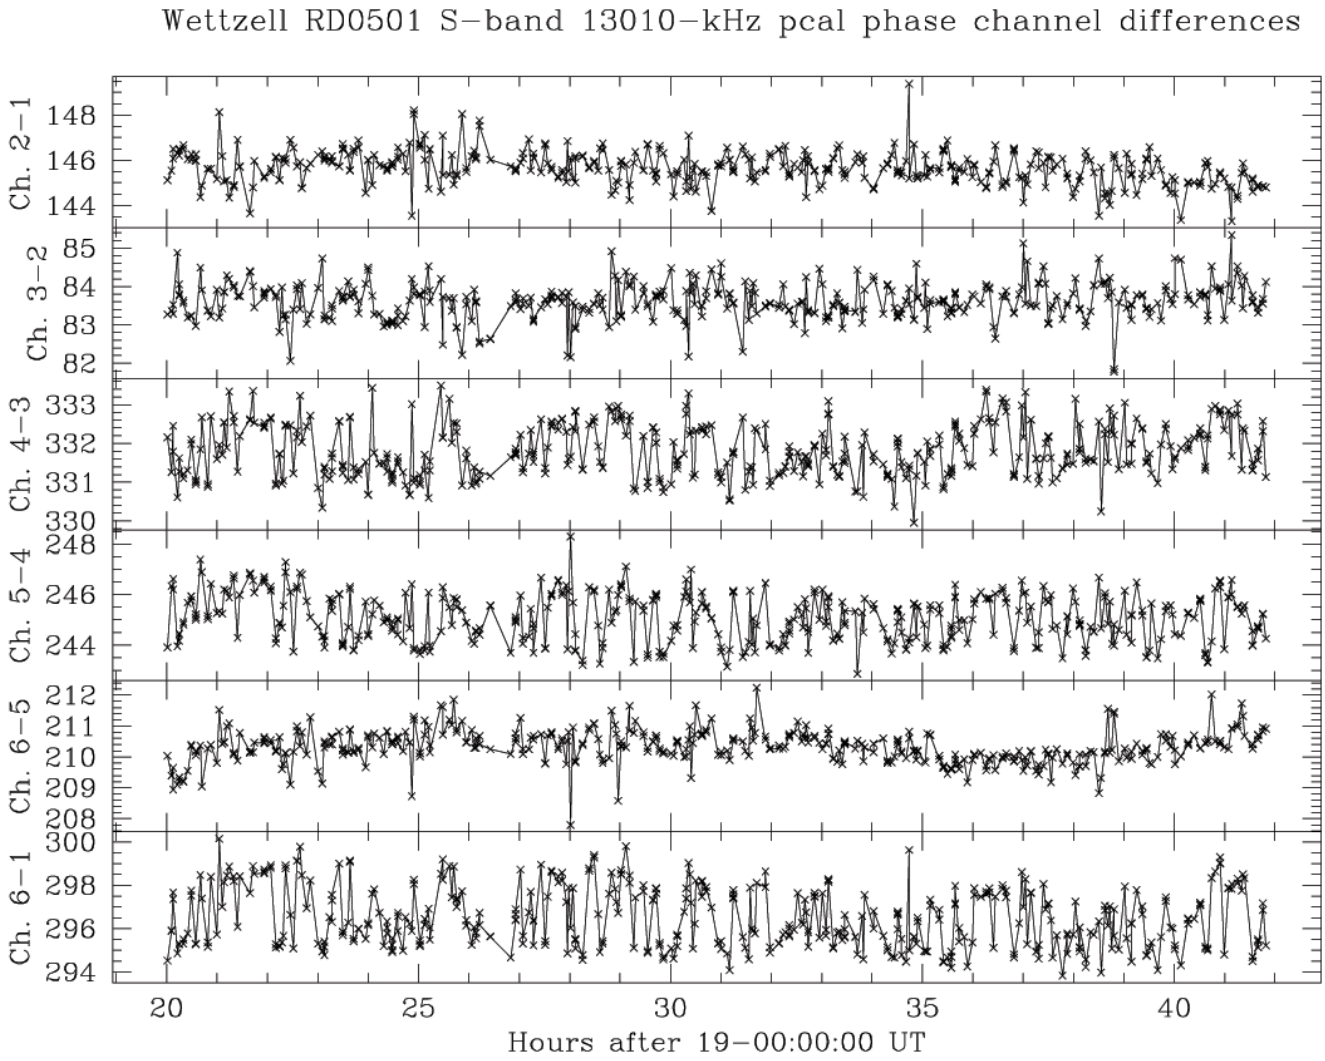

Note the lack of a systematic shift in the channel 4-3 or 5-4 phase cal phases during the hour-long interval of RFI in channel 4.

David B. Shaffer: RFI: Effects on Bandwidth Synthesis, IVS 2000 General Meeting Proceedings, p.402–406
http://ivscc.gsfc.nasa.gov/publications/gm2000/shaffer

RFI: Effects on Bandwidth Synthesis

David B. Shaffer

RadioMetrics/NVI, Inc./NASA Goddard Space Flight Center
e-mail: dbs@gemini.gsfc.nasa.gov

Abstract

I describe briefly how bandwidth synthesis works for VLBI, and then show how specific levels of RFI, expressed as a percentage increase above the nominal SEFD, affect the measured group delay.

1. Introduction

This note is a result of discussions held during an RFI workshop which took place at the Wettzell site on February 24, 2000 as part of the IVS meeting.

In the previous contribution, Brian Corey described radio frequency interference (RFI) measurements and how to determine the RFI power that may be incident on an antenna used for VLBI measurements.

The intent here is to describe quantitatively how RFI affects VLBI geodesy measurements and to quantify the level of RFI which is harmful.

2. Group Delay and Bandwidth Synthesis (BWS)

The most important quantity determined for a VLBI geodesy scan is the "group delay". This delay is the relative difference in the time of arrival of the random noise "signal" from a radio source at two antennas. As the Earth turns, the delay changes continuously, in a manner which depends on the relative location of the antennas and the position of the radio source. (Although many antennas participate in most VLBI experiments, in Mk III/Mk IV data analysis, the data from each pair of antennas is treated independently.)

Our ability to determine an accurate value for the group delay (at S- or X-band) depends on how much bandwidth is analyzed. Consider a wide bandwidth "signal" which is basically random noise. (This is what comes from the quasar radio sources.) The signal results from the addition of all the random electromagnetic fields within the overall bandwidth. This superposition of signals changes significantly on a time scale which depends on the maximum frequency difference (that is, the bandwidth). When we compare the two versions of this signal sampled (recorded) at different sites (this is what the correlator does!), we will find a non-zero correlation only if the two versions are closely aligned in time.

If B is the bandwidth that we process, the estimate of the group delay at which the signals are best aligned has an accuracy of about $1/B$. If the signal-to-noise ratio (SNR) is high, we can tell more accurately when the signals are aligned. In fact, the actual error (σ) in determining the best alignment delay is $\sigma = 1/(2\pi \cdot B_{rms} \cdot SNR)$. B_{rms} is the root-mean-square spanned bandwidth, which is about 40% of the total frequency span for the frequency sequences that we use.

How accurately must we measure the group delay?, *i.e.* what are the requirements on σ? We are trying to measure geodetic properties with a precision of better than a centimeter. Light

travels one centimeter in 33 trillionths of a second: $33 \cdot 10^{-12}$ seconds. One trillionth of a second is called a picosecond. Thus, we need to make group delay measurements with an accuracy of tens of picoseconds in order to achieve our desired geodetic goals. If we can achieve a modest SNR of 20, then we must analyze signals with bandwidths of hundreds of megahertz.

If we could record a complete bandwidth of several hundred MHz, our measurement of group delay would be relatively simple. We would just adjust the relative timing of the tapes until the correlator produced the maximum cross-correlation signal. However, the large bandwidth would require a very high data recording rate: twice the maximum bandwidth, or a sample rate approaching a gigabit per second. Although this rate should be achievable with the Mk IV, it is much higher than previous VLBI systems could record. Hence, a better way to sample a broad bandwidth was needed. Haystack's Alan Rogers showed in a 1970 paper that you don't need to record the entire bandwidth ("Very Long Baseline Interferometry with Large Effective Bandwidth for Phase Delay Measurements," *Radio Science*, **5**, 1239-1247). He demonstrated how to achieve nearly the same result by recording several narrow frequency channels spread out across the desired band. This technique, called "Bandwidth Synthesis" (BWS), is how we now record geodetic VLBI data, and is the reason for all the video converters in the data acquisition systems.

We must be careful how we combine the data from the BWS channels. In the case of a single very broad bandwidth, when the tapes are correlated at exactly the right delay, all parts of the processed bandwidth are in phase, with zero relative offset. If the estimated delay is not quite right, there will be a linear phase shift across the bandwidth. This shift arises because the real signal travel time between the two antennas does not match the delay picked by the correlator. In effect, the signal went a little farther (or a little less far) than the correlator estimated. At the higher frequencies in the band, the signal wavelength is shorter. Thus, at the higher frequencies, the extra distance that the signal went corresponds to more cycles (more phase) of the wavelength. If we can measure the shift of phase with frequency (the "phase slope") across the bandwidth, we can determine the difference between the delay used by the correlator (the "model delay") and the true delay. By adding the delay corresponding to the phase slope to the model delay, we get the actual delay. A set of such delays, determined for many scans, is then used to determine the overall geometry of the observations.

The BWS process takes advantage of this phase slope to measure the group delay correction to the correlator model: we process all the frequency channels, look at how the phase changes from channel to channel, and use this change to determine the delay.

This means we have to know how to align the various frequency channels. That is the purpose of the phase calibration system. The phase cal signals are injected in phase at the feed of the receiver. Any channel-to-channel variation of the phase cal phases when extracted from the individual frequency channels is due to differing path lengths through the VLBI equipment and differing video converter LO phases. The measured calibration phases are applied to the measured signal phases to take out the channel offsets. If the phase calibration fails, we can also observe a strong source to determine directly the phase offsets and apply the relative offsets to subsequent data: so-called "manual phase cal." This latter technique, however, does not let us keep track of time-dependent instrumental delay changes, so we really do prefer that the phase cal work!

Figure 1 shows a set of (fake!) BWS data, for our normal X-band observation mode, where eight frequency channels are used. The spacing of the channels is carefully chosen to maximize the effective spanned bandwidth and minimize confusing effects from combining the channels. At X-band, the channels that we use are often arranged in the sequence 0-1-4-10-21-29-34-36. These

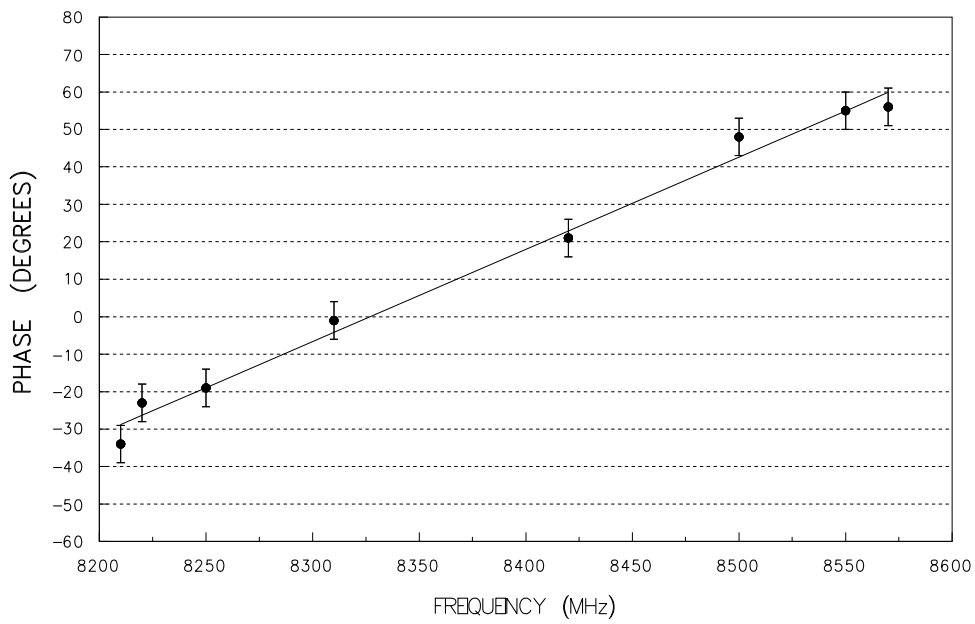

Figure 1. A (hypothetical) plot of phase versus frequency for an X-band experiment using the CDP standard narrowband sequence.

values denote the relative spacing between the channels. This sequence is usually multiplied by either 10 MHz or 20 MHz.

Because our equipment is not perfect and the phase calibration is not perfect, either, the amplitude of the correlated signal in the various channels is generally not the same, and there are residual phase offsets in each channel, too, as shown in Figure 1. Typical phase offsets are several degrees, resulting from such problems as phase offsets in the feeds, reflections in the phase cal system, and spurious phase cal signals. The amplitude and phase imperfections cause a group delay offset. As long as the phase offsets and the relative signal strength in each channel do not change, the group delay offset remains constant, and does not affect geodetic determinations. The offset is included in the clock offset term determined for each baseline.

3. RFI Effects

If the channel offsets change in a random sense during an experiment, there *will be* additional random group delay errors. If there are systematic changes, the delays will also be affected systematically. This is how RFI can cause serious problems.

Consider Figure 1. Determining the group delay is equivalent to fitting a line through the (frequency, phase) values for all the points. The slope of that line gives the group delay. Obviously, if the phase changes, a different slope (different delay) results. Generally, RFI does not cause phases to change. (However, some kinds of RFI—coherent signals at the phase cal frequencies—can cause the phase calibration phase to change, which will cause corresponding phase errors when the calibration is applied.) There is an error, however, on the measurement of the phase in each BWS channel. This error depends directly on the SNR in the channel. The size of this phase error *does* affect the fitting of the line through the phases: if the phase error in some channel is larger

(because the SNR is lower), that channel has a reduced effect on the fit. In the worst case, there may be so much RFI that a particular channel must be deleted from the fit. In this case, the effect of RFI is pretty obvious.

We can quantify the effect of RFI by noting how it affects the SNR in a frequency channel, and then how that affects the group delay determination. The channels for which RFI has the most effect are the end channels of the frequency sequences. These channels have the most leverage on the fit for the group delay. Since it is the overall SNR for the baseline (both antennas used together) that counts, we must relate the RFI level at one antenna to its effect on the interferometer.

The signal-to-noise ratio (SNR) of a VLBI observation depends on three parameters:

- the correlated flux density of the radio source: S_c
- the System Equivalent Flux Densities (SEFD) of the antennas
- the total number of bits correlated: N

Note that N depends on the scan length and the sample rate.

The exact expression for SNR is $SNR = S_c\sqrt{N}/\sqrt{SEFD_1 \cdot SEFD_2}$.

I now evaluate what happens to the group delay when a scan is affected by RFI at only one antenna. I assume that the correlated flux density, the SEFD at the second antenna, and the number of bits correlated do not change. Basically, the RFI increases the SEFD at the first antenna, which in turn reduces the overall SNR for the scan. The SEFD depends on the size of the antenna and the system noise power. (This power is usually expressed in terms of an equivalent noise temperature. Note that "temperature" and "power" are essentially equivalent, related by Boltzmann's equation: $P = kTB$.) When there is no RFI, the system noise temperature is determined by the internal noise power generated in the receiver amplifier, as well as the addition of some external radiation from the atmosphere and the ground. (Strictly speaking, the radio source itself also contributes to the system noise, but since most of the sources we observe add much less than one percent to the overall noise power, we usually neglect their noise contribution.) When there is RFI, the RFI power adds directly to the system noise power to increase the SEFD.

Thus, we can relate RFI to SEFD and then SNR by comparing the system power level with and without the RFI. If the RFI raises the power in a video convertor (as read by the Field System, using the `TPI` command, for example) by 10%, the effective SEFD in that channel will also increase by 10%, or a factor of 1.1, and the baseline SNR will be reduced by a factor of $1/\sqrt{1 \cdot 1.1} = 0.95$. If the RFI power is 100% of SEFD, it will double the power level, and the SEFD will also be a factor of 2 higher. Then the baseline SNR will be a factor of $1/\sqrt{1 \cdot 2} = 0.707$ lower.

Tables 1 and 2 show the effect RFI can have on group delay measurements when a single channel has a phase offset of 5 degrees, an offset which is not at all unusual. The RFI levels at one antenna are expressed as a fractional addition to the nominal system power level *in the affected channel only.* The "Relative Baseline SNR" shows the reduction in SNR, but only for the affected channel. I have applied the offset to the end channel of the frequency sequence to show the maximum effect. These tables show the group delay offset compared to the case of no phase error in the affected channel. The effect of RFI is to *reduce* the offset. This reduction, though, is not good—the original offset, without RFI, would just be absorbed into a clock offset. With RFI, what we consider in our analysis to be a clock offset is not constant!

The effect of RFI seems much worse at S-band. This is because the total frequency span is not as wide. In this case, a given phase offset has a larger effect on the slope of the phase line. Fortunately, the size of S-band delay errors is reduced by a factor of about 13 in our data analysis when the ionosphere correction is applied. (This correction combines the delays measured at S-

and X-band. We observe at two frequencies just so we can make this correction.)

Table 1. Effects of RFI at X-band

Single Antenna RFI Level	Relative Baseline SNR	Group Delay Offset
No RFI	1.000	16.9 picosec
10% RFI	0.953	15.9
20% RFI	0.913	15.0
30% RFI	0.977	14.2
40% RFI	0.845	13.4
50% RFI	0.816	12.8
100% RFI	0.707	10.3

Frequency Sequence is 0-1-4-10-21-29-34-36, multiplied by 10 MHz
RFI and a 5° phase offset occur in channel 8 only (frequency spacing 360 MHz)

Table 2. Effects of RFI at S-band

Single Antenna RFI Level	Relative Baseline SNR	Group Delay Offset
No RFI	1.000	96.9 picosec
10% RFI	0.953	92.1
20% RFI	0.913	87.9
30% RFI	0.977	83.9
40% RFI	0.845	80.3
50% RFI	0.816	77.1
100% RFI	0.707	64.0

Frequency Sequence is 0-1-4-10-15-17, multiplied by 5 MHz
RFI and a 5° phase offset occur in channel 6 only (frequency spacing 85 MHz)

At X-band, moderate RFI levels (those which increase the system noise power in one channel by less than 50%) can easily cause delay errors of several picoseconds, or more than 1 mm of geometric error. Larger RFI levels, of course, result in even larger errors, compared to the situation when there is no RFI. Notice that varying RFI (the usual case - RFI is seldom constant!) really causes a varying delay bias, rather than delay noise which could be either positive or negative. Thus, RFI will tend to "pull" the geodetic results, rather than just making the results noisier. For instance, if RFI is worse in a particular direction (a typical situation), the delays measured when the antenna is pointed in that direction will be affected systematically, leading to a biased position for the antenna. This biasing is probably the most serious reason why RFI is undesirable.

Based on simulations such as those used to generate the Tables, we have chosen an RFI level which causes a 10% increase in system noise power in a video converter bandwidth to be the level at which we begin to worry about degradation of VLBI observations. This is a quantifiable standard which we can present to other agencies as well as entities which are potential generators of RFI.

Brian E. Corey: RFI: Measurement Techniques, IVS 2000 General Meeting Proceedings, p.397–401
http://ivscc.gsfc.nasa.gov/publications/gm2000/corey2

RFI: Measurement Techniques

Brian E. Corey

MIT Haystack Observatory
e-mail: bcorey@haystack.mit.edu

Abstract

Some techniques for detecting radio frequency interference and for measuring its effect on the system temperature of a VLBI antenna are described.

1. Introduction

On 24 February 2000 a workshop on radio frequency interference (RFI) was held at the Fundamentalstation Wettzell during the IVS General Meeting. The topics discussed included (1) how RFI affects the VLBI observables (SNR, delay, etc.), (2) at what level RFI causes problems in geodetic VLBI, and (3) how to detect and measure RFI. The first two topics are treated in David Shaffer's companion paper, while the third topic is the subject of these notes.

Two general categories of measurement techniques may be distinguished by the type of equipment used:

- Measurements with a simple test set-up involving a small receiving antenna, an amplifier, and a spectrum analyzer.
- Measurements at a fully equipped VLBI station employing a VLBI antenna with receiver, a VLBI terminal, and a spectrum analyzer.

The first technique may of course be employed at a site with an existing VLBI antenna, but its most useful role is in surveying potential locations for future installation of a VLBI station. Specific examples of measurement procedures of both types are described in the next two sections.

As pointed out in Shaffer's paper, an increase as small as 10% in the power in a single baseband channel due to RFI can have a significant effect on the measured group delay. The methods described here are designed to be able to measure RFI-induced T_{sys} increases down to this level.

Most RFI sources are terrestrial and so are liable to be located near the horizon. RFI searches are therefore usually conducted by scanning the antenna around the horizon at, or slightly above, the minimum elevation angle normally observed.

2. RFI Measurements at a Site without a VLBI Antenna

Suppose we wish to characterize the RFI environment at a site where there will be a large VLBI antenna with antenna gain G_{vlbi} and system temperature T_{vlbi}. The goal is to detect RFI sources that would cause a 10% increase in the VLBI system temperature averaged over a bandwidth B_{vlbi} equal to the VLBI baseband bandwidth (typically 2 MHz or larger) anywhere in a given RF range, when the VLBI antenna is pointed more than a few beamwidths away from the source. (Ideally the search would be sensitive enough to detect sources so weak that they would increase T_{sys} by 10% only when the large antenna is pointed right at the source. Unless the RFI is much narrower

in bandwidth than B_{vlbi}, however, the procedure outlined here would require a search antenna as large as the VLBI antenna in order to detect such sources.)

The basic principle behind this method is the fact that, when a directional antenna is pointed more than a few beamwidths away from a source, the power received from the source is roughly the same as would be received by an omnidirectional antenna. There may be exceptions at some angles where the level is higher than this by (say) 10 dB, as when a prime-focus feed directly "sees" a source just beyond the edge of the primary reflector. But in general the power received away from the main beam will be within a factor of 2–3 of the omnidirectional level, or less.

2.1. Equipment and Configuration

- Test antenna with gain G_{ant} ($= 4\pi$ / beam solid angle in steradians, in the absence of losses in the antenna).
- RF amplifier with gain G_{amp}, with input attached directly to the antenna output.
- Spectrum analyzer that covers the desired RF range.
- Low-loss RF cable with loss L_{cab}, running from the amplifier output to the analyzer input.

See Figure 1. For convenience, define $G_{ac} = G_{amp}/L_{cab}$ [or, in dB units, $G_{ac}(\mathrm{dB}) = G_{amp}(\mathrm{dB}) - L_{cab}(\mathrm{dB})$]. G_{ac} is the net gain from the antenna output to the spectrum analyzer input.

2.2. Procedure

1. Set up the analyzer with frequency span covering the VLBI observing band, minimum input attenuation, narrow video BW (to minimize measurement noise – 10 kHz is often a good choice), and resolution BW B_{res} less than or equal to B_{vlbi}.
2. Scan the test antenna around the horizon slowly while observing the analyzer display.
3. Measure the total power P_{rfi} of each RFI signal observed. Doing this properly requires attention to the observed bandwidth B_{rfi} of the RFI signal. If $B_{rfi} = B_{res}$ (as will occur if the true signal BW is less than B_{res}), set P_{rfi} to the peak signal level. If $B_{rfi} > B_{res}$, set
   $$P_{rfi}(\mathrm{dBm}) = P_{ave}(\mathrm{dBm}) + 10\log\left[B_{rfi}(\mathrm{Hz})/B_{res}(\mathrm{Hz})\right] ,$$
   where P_{ave} is the average power level over the RFI bandwidth. If $B_{rfi} > B_{vlbi}$, set
   $$P_{rfi}(\mathrm{dBm}) = P_{ave}(\mathrm{dBm}) + 10\log\left[B_{vlbi}(\mathrm{Hz})/B_{res}(\mathrm{Hz})\right] .$$
4. Calculate the effective power $P_{omni} = P_{rfi}/(G_{ant}G_{ac})$ that would be received by an omnidirectional antenna, and the corresponding noise temperature averaged over bandwidth B_{vlbi}.
   Calculate power in dBm:
   $$P_{omni}(\mathrm{dBm}) = P_{rfi}(\mathrm{dBm}) - G_{ant}(\mathrm{dB}) - G_{ac}(\mathrm{dB})$$
   Convert to milliwatts:
   $$P_{omni}(\mathrm{mW}) = 10^{[P_{omni}(\mathrm{dBm})/10]}$$
   Convert to noise temperature:
   $$T_{omni}(\mathrm{K}) = P_{omni}(\mathrm{mW}) \,/\, [1.4 \times 10^{-20} \times B_{vlbi}(\mathrm{Hz})]$$
   T_{omni} is the typical noise temperature increase due to RFI that will occur in a baseband channel when the VLBI antenna points away from the RFI source. If the VLBI antenna points directly at the source, the increase will be approximately $T_{omni}G_{vlbi}$, unless the source is in the near field (see Section 3), in which case the increase will be smaller than this.

5. Repeat steps 2-4 with other analyzer resolution bandwidths narrower than B_{vlbi}. The narrower B_{res} is, the greater the sensitivity to narrowband RFI, but at a cost of longer analyzer sweep times.

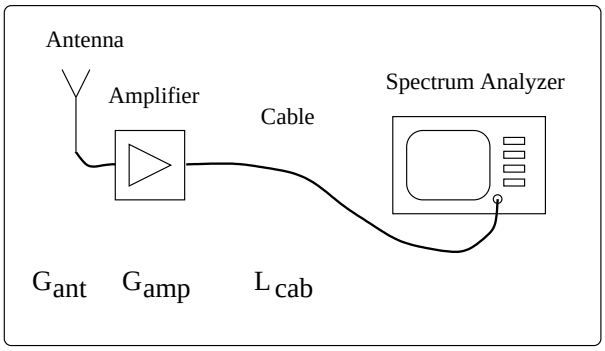

Figure 1. RFI measurement set-up without a VLBI system (courtesy of Armin Böer).

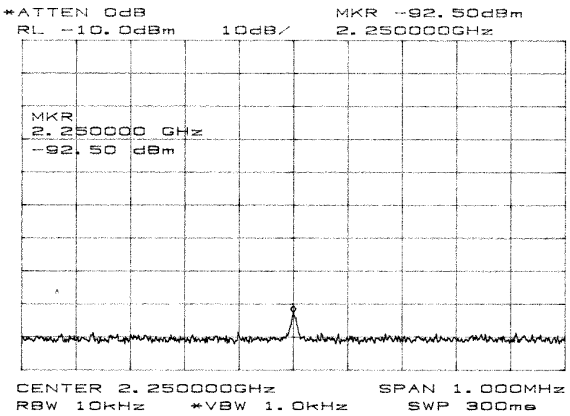

Figure 2. Sample analyzer display showing RFI (courtesy of Gerhard Kronschnabl).

2.3. Example

Figure 2 shows an RFI spike observed with the following equipment: an S-band horn with gain G_{ant} of 14 dB, an Avantek ABG240 amplifier with gain G_{amp} of 22 dB, an RF cable with loss L_{cab} of 6 dB, and an HP 8561A spectrum analyzer. The combined amplifier+cable gain G_{ac} is $22 - 6 = 16$ dB. From the figure, the observed signal bandwidth is the same as the resolution bandwidth, and P_{rfi} is equal to –92.5 dBm. Step 4 above then gives:

$$\begin{aligned} P_{omni} &= -92.5 \text{ dBm} - 14 \text{ dB} - 16 \text{ dB} \\ &= -122.5 \text{ dBm} \\ P_{omni}(\text{mW}) &= 10^{-122.5/10} = 0.56 \times 10^{-12} \text{mW} \end{aligned}$$

For a baseband bandwidth B_{vlbi} of 2 MHz, the effective temperature would then be

$$\begin{aligned} T_{omni}(\text{K}) &= 0.56 \times 10^{-12} \text{ mW} / [1.4 \times 10^{-20} \times 2 \times 10^{6} \text{ Hz}] \\ &= 20 \text{ K} \end{aligned}$$

2.4. Requirements on the Equipment

Assume that the weakest RFI source that can be detected reliably is 3 dB above the noise floor P_{nf} on the analyzer, so the minimum detectable level is P_{nf}. Let T_{nf} be the equivalent noise temperature given by

$$T_{nf}(\text{K}) = P_{nf}(\text{mW}) / [1.4 \times 10^{-20} \times B_{res}(\text{Hz})]$$

T_{nf} for spectrum analyzers is typically 100,000 K or higher. Therefore, in order to be able to detect sources that could interfere with systems having noise temperatures below 50 K, it is essential to have an antenna and an amplifier with high gain and low noise ahead of the analyzer.

To be quantitative: Let T_{test} be the noise temperature of the test antenna and amplifier together, in the absence of RFI; T_{test} will include contributions from amplifier noise, atmospheric noise, and especially ground pickup. The noise floor on the spectrum analyzer is set by the sum of the antenna+amplifier noise at the analyzer input and the internal analyzer noise $T_{nf,sa}$:

$$T_{nf} = T_{test} G_{ac} + T_{nf,sa}$$

The requirement that the minimum detectable T_{omni} be less than 10% of T_{vlbi} is then equivalent to

$$(T_{test} G_{ac} + T_{nf,sa})/(G_{ant} G_{ac}) < T_{vlbi}/10$$

or

$$T_{test}/G_{ant} + T_{nf,sa}/(G_{ant} G_{ac}) < T_{vlbi}/10$$

In a well-designed system, the total system noise will be dominated by the noise from the antenna+amplifier, not by the internal spectrum analyzer noise. This requirement together with the last equation then yield

$$G_{ant} > 10\, T_{test}/T_{vlbi}$$

and

$$G_{ac} > T_{nf,sa}/T_{test}$$

For T_{test} = 300 K, T_{vlbi} = 30 K, and $T_{nf,sa} = 10^6$ K, these requirements become $G_{ant} > 100$ (= 20 dB) and $G_{ac} > 3000$ (= 35 dB).

3. RFI Measurements with a VLBI Antenna

Using a large antenna to look for RFI has the advantage of high sensitivity. A disadvantage is the narrow beamwidth, which necessitates scanning the antenna slowly if maximum sensitivity to RFI is desired. For example, a 20-meter antenna has a beamwidth of ~0.13° at X-band; a full sweep around the horizon at a single elevation angle with an integration time of 1 second per beamwidth would therefore take nearly an hour. [Note that, if a source is in the "near field" of the antenna, which extends out to distances of 2×(antenna diameter)^2/wavelength, or ~20 km for a 20-meter antenna at X-band, the effective antenna beamwidth is larger than the standard "far field" value measured on celestial sources. In other words, nearby sources appear larger.]

3.1. Measurements Using Baseband Detectors

A simple method for detecting RFI uses the baseband total power detectors built into most VLBI terminals. For example, MkIII/MkIV video converters and VLBA baseband converters all have total power integrators whose output can be read via the PC Field System command `tpi`. A procedure utilizing the `tpi` command is:

1. Set up the VLBI terminal as for an experiment, including frequencies and bandwidths.
2. Use the `tpi` command to log the baseband power levels in all channels every 1-2 seconds. For example, an appropriate command for a MkIII/MkIV system is
   `tpi=v1,v2,v3,v4,v5,v6,v7,v8,v9,v10,v11,v12,v13,v14@!,1s`
3. Scan the antenna around the horizon no faster than one beamwidth per `tpi` readout interval.

4. After the scan, examine the data for variations larger than 10% from sample to sample and between frequency channels. There may be some variability with time in the power levels due to weather (try to observe during good weather conditions!) and due to changing ground pickup (the horizon is unlikely to be perfectly flat!). Variability due to these two causes can usually be distinguished from RFI by the fact that clouds and ground generally increase T_{sys} by the same amount at all frequencies within a band (i.e., S or X), whereas RFI generally affects only a few channels at most.

3.2. Measurements Using a Spectrum Analyzer

A spectrum analyzer may be used instead of, or in addition to, the baseband power detectors to search for RFI. If the RFI is weak (raising T_{sys} by less than a factor of 2) and broadband (on the order of the bandwidth of a baseband channel or wider), the baseband detectors generally provide a more sensitive search method. But if the RFI is strong or narrowband, the spectrum analyzer approach works well and has the advantages of instant visual feedback and full IF coverage. (In the baseband detector approach, the measurements must be repeated for each observing frequency sequence.) To search for RFI with a spectrum analyzer:

1. Connect an IF signal from the receiver to the spectrum analyzer. Set the analyzer for a frequency span covering the full IF range, a narrow video BW, and a resolution BW B_{res} no larger than the baseband BW B_{vlbi}. Be sure the signal observed on the analyzer display is dominated by the receiver IF and not by internal analyzer noise.
2. Scan the antenna around the horizon at a rate no faster than one beamwidth per analyzer sweep time. If the sweep time is too long (and hence the scan rate is slower than desired), the video BW may be increased, but at the cost of increased noise.
3. For each RFI signal observed on the analyzer, note the signal bandwidth B_{rfi}. If $B_{rfi} = B_{res}$, set P_{rfi} to the peak signal level; if $B_{rfi} > B_{res}$, set P_{rfi} to the average power level of the RFI. Also note the power level P_{nf} of the noise floor away from the RFI.
4. Estimate the fractional increase $f = T_{rfi}/T_{sys}$ in T_{sys} due to the RFI signal according to

$$\begin{aligned} f(\mathrm{dB}) &= P_{rfi}(\mathrm{dBm}) - P_{nf}(\mathrm{dBm}) + 10\log\left[B_{rfi}(\mathrm{Hz})/B_{vlbi}(\mathrm{Hz})\right] \\ f(\mathrm{numeric}) &= 10^{f(\mathrm{dB})/10} \end{aligned}$$

   The noise temperature due to the RFI is then $f \times T_{sys}$.
5. Repeat steps 2-4 as desired with narrower B_{res} settings in order to detect narrowband RFI at weaker levels.

Example: Suppose the plot in Figure 2 was obtained as described in this section. Then $P_{rfi} = -92.5$ dBm, $P_{nf} = -101$ dBm, and $B_{rfi} = 10$ kHz. By step 4 above, this signal would increase the system power in a 2 MHz bandwidth by the fractional amount

$$\begin{aligned} f(\mathrm{dB}) &= -92.5\ \mathrm{dBm} - (-101\ \mathrm{dBm}) + 10\log\left[10^4\ \mathrm{Hz}\ /\ 2\times10^6\ \mathrm{Hz}\right] \\ &= -14.5\ \mathrm{dB}\ , \end{aligned}$$

which equals 4%.