

RFI and How to Deal with It

J. R. Fisher

NRAO, P. O. Box 2, Green Bank, WV 24944, rfisher@nrao.edu

Abstract. Interference from active users of the radio spectrum will always be a problem for radio astronomy, but careful planning, active data monitoring, and a few data editing techniques allow us to use a more spectrum for science than one might expect. A good plan starts with knowing what to expect in the frequency range that you want to use. Then there are a number of receiver options, signal processing configurations, and observing strategies that will help to minimize the effects of interference to your data. The degree to which you can remove interference from data after it is recorded depends considerably on having chosen the best observing parameters before the observing run begins. Research into real-time signal processing for interference excision offers some hope that future radio astronomy receivers will better separate cosmic signals from man-made transmissions. Some of these excision techniques under investigation are reviewed.

1. Introduction

Many radio astronomers are accustomed to observing in the protected radio astronomy frequency bands or at high enough frequencies where radio frequency interference (RFI) is, at most, a minor annoyance. To a large extent this will continue to be the case, but the advent of wider band receivers and spectrometers and an increasing desire to pursue science outside of the protected bands present new challenges. A few frequency ranges are so heavily filled with radio transmitters that their use for astronomy is unrealistic. Other frequency bands are sparsely filled or used intermittently on time scales from microseconds to days. If the science is important enough, a careful study of the spectrum and RFI signal characteristics may permit observations that, at first glance, would appear impractical. Even where interference is isolated in frequency or is present only a small fraction of time, a little advanced planning can decrease the fraction of data that must be discarded due to RFI contamination.

To focus this discussion I will concentrate on spectral line observing at Arecibo and Green Bank. Many any of the same considerations apply to continuum and pulsar observations and to other observatories. Continuum observers may want to consider the use of a spectrometer for their observations at frequencies where RFI is a problem since this offers the possibility of rejecting interference in the frequency domain. Data rates and data volumes made this impractical in the past, but faster processors and much larger storage capacities have changed the picture considerably.

Much of what I have to say will be self-evident, but it can serve as a check list for setting up an observing run when you know that RFI could be a problem. I encourage you to develop your own techniques and algorithms for interference removal and contribute these to the community.

2. Early Planning

Several resources are available at Arecibo and Green Bank for getting a preview of the interference environment that you will encounter. These are of most help in writing your observing proposal and in communicating with the observatory staff before arriving for your observing run. If appropriate, learn the name of the contact for interference matter at the observatory, and make them aware of your observing plans and the frequencies that you expect to use to give them a chance offer advice on recent RFI experience and, in a few special cases, to coordinate your observing times with users of the spectrum with whom a coordination agreement exists. Don't rely on a comment on your observing proposal to catch the eye of the local spectrum manager.

Both Arecibo and Green Bank have web pages devoted to interference matters that may be accessed from their home pages at the following addresses:

<http://www.naic.edu> click [Scientific Users] [Spectrum Management]

<http://www.gb.nrao.edu> click [Quiet Zone] and [RFI Protection]

Spend a little time browsing these web pages to familiarize yourself with the RFI activities at the respective sites.

2.1. Spectrum surveys

The Arecibo Observatory has an antenna and spectrum analyzer on the hilltop above the telescope near the lab. This antenna provides a recent survey of the spectrum between 0 and 1450 MHz where the interference is most dense. These measurements may be found at

<http://www.naic.edu/rfuser/smarg-hplots.html>

There are also a number of statistical plots and lists of frequencies of known signals seen in the two 20-cm (L-band) receiver frequency ranges which can be found at

<http://www.naic.edu/rfuser/smarg-vig1.html>

At Green Bank spectrum surveys are available for the frequency ranges of most of the receivers that were available on the 140-foot telescope between 50 MHz and 3.4 GHz. These surveys were done with the FFT spectrometer at the receiver output with the telescope pointed at the zenith. The spectral resolution was 40 kHz. The spectrum plots and an explanation of the data acquisition and analysis may be found at

<http://www.gb.nrao.edu/RFI/>

A new spectrum survey of GBT receivers is in the works.

2.2. Coordination

Spectrum sharing is becoming a more frequently used spectrum management tool, particularly in the allocated bands where radio astronomy is a co-primary or secondary “service” in the band as defined by the International Telecommunications Union. At least one sharing agreement is already in place with the Iridium satellite operators and others may be negotiated in the future. These agreements usually require some sort of expected use forecasts from the observatories for scheduling purposes. Also, the Arecibo Observatory has established informal coordination procedures with users of the radio spectrum near the most frequently used observing bands.

To make these coordination procedures work, information about your expected frequency use during your observing run should be provided to the observatory’s spectrum management staff well in advance of your observing run. At the present time there are very few observing programs that will be affected by coordination procedures, but it is a good idea to check the latest status of such agreements. The most frequently used coordinated band is at Arecibo below 1400 MHz where radar operators may be asked to limit their transmissions to their lowest frequencies when given sufficient notice. See the web page

<http://www.naic.edu/~rfiuser/smarg-coord.html>

for a description of current agreements.

2.3. Interference characteristics

The published spectrum survey, where available, will tell you roughly what to expect in the way of RFI in your observing band. If interference looks like it is going to be a problem, you will want more detail about spectrum usage. Requests to the observatory for more detailed measurements of particular parts of the radio spectrum are welcome with the caveats that the appropriate equipment must be available and that the requests can be accommodated within the work load of the staff. This is an area where the dialogue between observers and spectrum managers can be enhanced with better communication. Well specified and limited requests are the most likely to be within the realm of possibility.

Both observatories are working to provide better information on current spectrum use. A useful data set would be something like a time-frequency matrix of the telescope receiver output for the times of day when we expect to make our observations. A time resolution of one second and a frequency resolution of 1 kHz seem like good starting points. Eight hours of floating point data in a 10 MHz wide spectrum would require 1.2 gigabytes of data storage, which isn’t unrealistic. From this we could determine the fraction of time which is likely to be available for astronomical measurements at each frequency in the spectrum and try summing portions of the data to look for very weak interference. With a little more work we might identify many of the signals and infer a bit more about the usage of this spectrum.

Of course, it is unlikely that we can ask for telescope time for a spectrum sample of this type. A close approximation to the same sensitivity can be obtained with a well placed RFI sampling antenna and a reasonably low noise preamplifier. The typical gain of radio telescope sidelobes to RFI is about -10 dBi so a sampling antenna with 0 dBi gain and ten times the system temperature

is roughly equivalent. The other half of the sensitivity equation is integration time so the RFI sampling receiver must be connected to an integrating back-end, such as a Fourier transform spectrometer or autocorrelator, rather than a swept-frequency spectrum analyzer. Only observer experience and feedback can tell whether there is enough the benefit to the telescope's science to warrant the allocation to such resources.

Armed with the ideal spectrum information one can then ask the question, which of the RFI signals are from licensed users of the radio spectrum and which are from incidental emissions from equipment in the vicinity of the observatory. If the receiving antenna is fixed, and a signal does not change much in intensity over time, there's a good chance that it can be tracked down and turned off. This detective work can be time consuming, but it has general benefits to the observatory's environment.

3. Observing Considerations

With a good picture of the interference environment you can plan your observing strategies accordingly. Some questions to begin answering are: What dynamic range is required of the spectrometer? What total bandwidth is feasible and what frequency resolution is required to isolate the interference? What time resolution is necessary to take advantage of interference-free intervals? Are there any very strong signals that need to be avoided?

3.1. Spectrometer type and configuration

Some spectrometers are more robust to RFI signals than others. The two-bit sampling autocorrelators at Arecibo and Green Bank and the six-bit sampling FFT spectrometer at Green Bank are the extreme examples. An interfering signal whose power is a significant fraction of the total noise power in the spectrum will distort the output of a two-bit sampler and cause problems to the spectrum even well separated in frequency from the offending signal. The six-bit FFT spectrometer can swallow a signal ten times the total noise power with relatively little spectral distortion well away from the strong signal. Both the Arecibo and Green Bank autocorrelators have a three-bit sampling mode, which has an intermediate dynamic range. The penalty for using the spectrometers with larger sampling words is that fewer spectral channels and/or narrower total bandwidths are available. You need to look at the spectrometer specifications to see what trade-off suits your needs best. At frequencies where RFI is a major consideration the widest spectrometer bandwidths are often not necessary, so you can usually find a useful compromise.

Autocorrelators have a reputation for “ringing” of narrowband interference than tends to destroy an entire spectrum that otherwise would be undistorted. By this we mean that the instrumental profile (called a point spread function by opticians) of a single channel of an autocorrelation spectrum is a $\sin(\pi x)/\pi x$ function, where x is the separation from the narrowband signal in number of channels. This looks like a damped sine wave on either side of the strong signal than can cover most of the spectrum. The ringing from weak to moderate signal strengths can be effectively suppressed with a “Hanning” convolution function. Better convolution functions are available to deal with stronger signals, but I

am not aware of any investigations on whether sampler output distortions or residual “ringing” is the limiting factor for strong signals. My guess is that the modern 3-bit correlators could benefit from better convolution functions, so some of the spectral “ringing” folklore is probably out of date.

In an FFT spectrometer the spectrum convolution must be done in real time. You do not have the option of changing your mind about the convolution after the data are recorded as you do with an autocorrelator’s spectrum. An FFT spectrometer does not do the convolution in the frequency domain because it is computationally cheaper to do the equivalent by tapering or windowing the voltage samples fed into the FFT in the time domain. The taper function must be selected before your observing begins. Because tapering underweights some of the data entering the FFT, some loss of sensitivity is associated with tapering. This is another trade-off to weigh in your observing setup.

At frequencies where interference is not a problem, spectrum integrations of 60 seconds or more before storing a data unit on disk are commonly used to keep the recorded data volume reasonably small. If, from a high-time-resolution spectrum sample, you find that shorter integrations are necessary to make use of RFI-free intervals, a much shorter integration time may be justified. On-the-fly mapping uses integration times as short as 100 milliseconds or so, so it is not completely unreasonable to do the same for RFI editing. There may not be a great deal to be gained from really short time resolutions so consider the merits carefully before being presented with a truckload of Exabyte tapes or CDs.

At some time scale less than a second or so it makes more sense to throw away contaminated data in real time. Several possibilities currently exist, although they are not in common use or have not been tried at all. At Arecibo Bill Sisk and colleagues have built a very effective blanking pulse generator that synchronizes itself with the 1330/1350 MHz FAA radar signal at Pico del Este (Hagen, 1988). When this pulse is applied to the autocorrelator it turns off the spectrum integrators for a selected time interval, typically 512 microseconds, to allow the radar pulse and echoes to clear out of the correlator delay line before resuming integration. The radar pulses are spaced about 3 milliseconds apart so not too much data is lost. Be aware, however, that in its current configuration the autocorrelator waits until it has integrated for the specified amount of time before recording it on disk so the clock time duration of a scan will be extended proportionately.

Another blanking possibility is that the Iridium satellite system provides a blanking pulse for the duration of 50% of the 90-millisecond on-off period of the transmission cycle. During the off fraction the satellites are not transmitting in the 1612-1625 MHz band so, in principle this time is available for radio astronomy. The useful interval available is smaller than 45 milliseconds because of the differences in delays to various satellites in the sky. The two correlators in Arecibo and Green Bank and the FFT spectrometer in Green Bank can accept this blanking pulse, but it is not commonly used so ample notification needs to be given to the staff to implement and test it.

Any other blanking pulse that can be generated in synchronism with burst type interference is a candidate for use with one of the spectrometers. At present no others are routinely provided, mainly because there has been no specific requirement. The FFT spectrometer at Green Bank has a built-in burst detector

that proved reasonably effective on lightning interference, but it was so hard to set up that it fell into disuse. It may be resurrected in a more modern form as priorities dictate.

3.2. Receiver filters

A radio astronomy receiver typically has a total amplifier gain of 100 dB or more. Hence, a strong interfering signal will overload the receiver somewhere in the amplifier chain if it is not first attenuated by a frequency-selective filter. A few signals, like the radar signals near 1300 MHz at Arecibo and Green Bank and a few satellite signals can overload receivers in the telescope far sidelobes after only 40 to 50 dB of receiver gain. These signals require filters immediately following the cooled low-noise amplifier stages. When a receiver overloads it is not useful anywhere in its passband.

Check the published frequency lists of known strong signals and consult with the observatory staff about the optimum filter and frequency conversion setup for the frequencies that you intend to use. A standard setup for the receiver that you intend to use probably takes into account the existence of known strong signals, but it's a good idea to understand the filter configuration that you will be using. If you don't have a block diagram of the receiver system, ask for one, or ask someone to draw one for you. Many observers take for granted that the receiver works and don't particularly care about the details, so we have fallen out of the habit of providing lots of system information. In demanding situations, such as a high RFI environment, you shouldn't be shy about asking to understand more about your observing setup. Most engineers will be pleased that you are interested.

Keep in mind that all mixers have two sidebands. If the sidebands are not very well separated in frequency, the unwanted (image) sideband may not be suppressed sufficiently to reject a strong signal in its frequency range. Single-sideband mixers are especially vulnerable to unwanted sideband interference because they are limited to about 30 dB of image rejection. The FFT spectrometer uses a single-sideband mixer in its conversion to baseband, so pay special attention to this one. Think about arranging an RF or IF filter to reject signals that would appear in the unwanted sideband. See the section on diagnostics below for a discussion on how to determine whether an interfering signal is leaking through an unwanted sideband.

3.3. Optical fibers

A subtle dynamic range issue has been introduced in to the receiver equation by the use of analog optical fibers for the transmission of wideband signals from the telescope receiver room to the control room. Present day optical fiber modems have a very high equivalent noise temperature on the order of 5×10^6 Kelvins so they require about 80 dB of gain ahead of them. A strong RFI signal can overload the amplifier stage immediately before the modem or the modem itself. Whenever possible, use the narrowest available filter before the input to the fiber modem that will pass your observing frequency band. As with all receiver stages, don't use any more gain ahead of the fiber modem than is necessary to achieve a good overall noise performance. We tend to concentrate on optimum receiver temperature and use a little more gain throughout the system than is

absolutely necessary, so there may be some room for dynamic range optimization in tight RFI situations.

3.4. Diagnostics and reporting

Once an observing run has begun there isn't much time to deal with interference problems because you will want to be concentrating on science issues. Unless the data are completely wiped out, we cannot expect an observer to spend much time diagnosing RFI. Nevertheless, the observatory staff would like to hear about interference during your observing run and to know as much about the observing setup and interference characteristics as possible. If you have time to run a few diagnostics, all the better. Sometimes the RFI is generated locally or even within the receiver itself, and there may be something that can be done about it with prompt notification.

The following information about observed RFI is particularly helpful:

Frequency - What is the exact measured sky frequency of the signal? At what frequency does this same signal appear in the IF(s) and/or baseband? What sideband are you using, if you have a choice? If you move the observing center frequency, does the signal remain fixed in sky frequency, intermediate frequency, or neither?

Time - When is the interference seen? What is its duration or other temporal characteristics? Is it pulsed, steady, modulated, or periodic?

Bandwidth - Is the signal unresolved in your spectrum or does it have frequency structure? Does it drift in frequency with time? Sample spectra of the interference with clearly marked frequency and intensity axes are very helpful.

Direction - Does the signal have any dependence on the pointing direction of the telescope?

Polarization - Is the signal stronger in one polarization than the other? Does this change with time, telescope position, or feed rotation?

Observing conditions - What integration time, switching mode, polarization, LO settings, and filter selections were you using? Are there any other parameters that might be relevant?

The purpose of this information is to isolate the source of the interference. If the RFI source is external to the telescope it will generally vary in intensity with time, telescope orientation, and polarization. If the signal enters the system ahead of the first mixer it will have a constant sky frequency but its intermediate frequency will change when the first LO frequency is changed. If the signal's intermediate frequency changes in the opposite direction from what you expect, it is entering through the unwanted sideband. If the RFI signal's intermediate frequency changes by more than or a strange multiple of the LO frequency change, the interfering signal may be the result of a spurious mixing product in the receiver system.

The signal's bandwidth is a good hint to its origin. Very narrowband signals at telltale frequencies, such as exact integer multiples of 1, 5, 10 or 100

MHz, usually originate somewhere in the observatory. Broadband signals without much frequency structure are often associated with lightning, relay arcing, welders, fluorescent lights, ignition noise, etc. This type of radiation can often be identified by its temporal structure. For example, heater thermostat relays typically cycle on a 3 to 20 minute period, and ignition noise has a pulse frequency associated with gasoline engine speed. Most broadband noise is stronger at lower frequencies, mostly below 1 GHz.

The spectrum of RFI from computers and other digital equipment is a mixture of narrow and broadband emissions with the narrowband spikes being spaced at intervals from kilohertz to many megahertz. There is usually some periodicity to the narrowband frequencies, but there will also be a jumble of spikes that don't fit any well-defined pattern. A common preconception is that computers radiate primarily at their CPU clock frequency, but the intensity at this frequency is often no higher than at a lot of other frequencies in the spectrum. Most digital RFI is strongest below a few hundred megahertz, but strong digital noise can be seen to several gigahertz and above.

The signals from active users of the radio spectrum come in a wide variety of bandwidths and temporal characteristics. These are beyond the scope of this lecture so I suggest that you contact the observatory's spectrum manager for information on the signals expected in your observing frequency range. The experience of local RFI experts can also be very helpful in identifying interfering signals.

4. Data Reduction

I think it is fair to say that, currently, there are very few tools available for RFI mitigation in the standard data analysis packages, mainly because there has been little demand. There are many reasons for this. No single or well-defined solutions to RFI excision exist. Automated excision has tended to be only marginally effective. Data rates and storage volumes have not been adequate for RFI isolation. Finally, observers tend to avoid RFI or be resigned to its existence. If the science warrants a greater effort on RFI mitigation, interested observers can gradually change this situation by being specific about the RFI problems that need to be solved and by contributing any tools for RFI excision at the data analysis level that they find to be effective.

4.1. Convolution/ACF tapering

As mentioned earlier, the native instrumental profile of a channel in an autocorrelation spectrum is $\sin(\pi x)/\pi x$. This function has quite significant amplitude many channels away from a narrowband signal, which means that a strong interfering signal can affect a large portion of the spectrum. This instrumental profile can be made to fall off much more rapidly by either convolving the spectrum with an appropriate function or tapering the autocorrelation function before performing the Fourier transform to produce the power spectrum. The most common convolution function, called a Hanning function, is a [0.25, 0.5, 0.25] weight of three adjacent channels. This is computationally easy to do and is quite sufficient for interference of modest intensity. More sophisticated suppression of the instrumental profile sidelobes is probably better done with tapers

to the autocorrelation function. A few useful taper functions may be found in Rabiner and Gold (1975). One or more of these should be added to the standard data analysis packages.

4.2. Manual editing

The most common form RFI excision in data analysis consists of manually deleting spectra that are obviously contaminated more extensively than average. If there aren't too many data records per scan, this can be done by simply examining each record and deleting contaminated spectra from the average of all spectra in each scan. Keep in mind that deleting a data record will change the effective time over which the data are averaged. If you are using total power on-off position switching, it is often a good idea to delete corresponding records in the on and off scan pair, even though only one of the two records is affected. This preserves the intrinsic symmetry of the on-off pair that is often important to accurately subtract baseline distortions due to changing ground or atmospheric noise input to the system.

A more informative and flexible way of editing data is to display the data from one or a pair of scans on a grey scale plot of frequency vs time, where each line in the plot is a spectrum from a single data record. Even if you have recorded a large number of data records per scan by selecting a short integration time, the data will generally fit nicely on a single image. Weak interference that is distributed over adjacent pixels will tend to stand out better in a grey-scale plot as will subtle large-scale changes in the spectra over the scan. RFI can be edited with a few graphical tools, such as drawing a box around data in the time-frequency domain or marking ranges of data records for deletion. If you delete an isolated domain in both time and frequency space, the frequency channels that have less data in the scan average will have a higher rms noise value. This will lead to artifacts in your spectrum that you must be careful not to interpret as real signals. Deleting all records at a given frequency isn't terribly useful since this just leaves a hole in the averaged spectrum.

A useful grey-scale display requires that the receiver gain be normalized across the spectrum. A good normalizing spectrum may be hard to derive if interfering signals are present for the whole scan since they will tend to get normalized out of the display. One can play a few tricks, such as averaging or median filtering the normalization spectrum in the frequency domain to accentuate narrowband signals in the grey-scale display.

With or without effective data editing, we'd like to have some measure of the quality of data that went into a scan average. Displays of the peak, rms, and statistical asymmetry of each spectral channel's data are a few quality measures that can be useful. Since most astronomical signals are weakly polarized, a comparison of intensities in the two receiver polarization channels can be a very useful flag for the presence of weak interference.

4.3. Algorithms

Automating the process of manually editing data to remove RFI is surprisingly difficult. It is hard to duplicate all of the factors that go into a judgement about which data are valid and which are contaminated. An editing algorithm requires one or more unambiguous properties that distinguish RFI from normal

data. The more noise-like an interfering signal looks in its intensity distribution, the more difficult it is to remove from data. An effective algorithm often needs more than a few built-in assumptions about the nature of the interference that it is intended to remove.

If an interfering signal is strong and either on or off, it can be detected and rejected with a threshold algorithm. Either a range of channels around the interference frequency or entire spectra may be discarded when the interfering signal is detected.

A more likely situation is that an RFI signal will vary continuously and randomly in intensity from strong to undetectable. A threshold detection algorithm will not discard weak interference which will usually appear quite prominent in the integrated spectrum for the full scan. If the interference has a fairly low duty cycle, say less than 30%, using a median instead of a mean average for each channel in the scan may be more effective. The noise in a median is not much worse than for a mean as long as the amplitude distribution for each spectral channel is very nearly gaussian. If there is a gain drift which is a significant fraction of the standard deviation of individual channel noise in one record, the median spectrum for a scan will be considerably noisier than a mean spectrum without interference. If the gain variations are uniform across the spectrum, each record's spectrum may be normalized to an average gain before computing the median for each channel.

Simply discarding spectral channels in the scan average that contain RFI is usually not very useful because the astronomical information in those channels is then lost. However, if you can afford to record spectra with much better spectral resolution than is required by the science, then very narrowband interference may be removed by replacing isolated channels with averages of adjacent channels. Since one's eye can ignore isolated channel interference, this type of rejection tends to be of a cosmetic nature.

5. Real-Time RFI Excision

Successful excision of interference with data editing relies on the fact that the interference is present for a reasonably small fraction of the time within the time resolution that one can afford to record for later processing. If the RFI is present most or all of the time at the frequencies of interest, then it must be isolated in the spatial domain or some combination of time and spatial domains that clearly distinguish it from the natural cosmic signals. A simple example would be to place an antenna null in the direction of the interfering signal so that the radio telescope doesn't see it.

Active nulling of interference has been used by the radar, communications, and acoustics industries for some time. Noise-cancelling headphones are a good example. While some of their basic techniques are applicable to radio astronomy, the degree to which interference must be rejected is quite different. Astronomical signals are almost invariably at the detection limit of our instruments after many minutes of signal integration. A weak interfering signal is nearly as destructive as a strong one and can even be more harmful because of its subtle nature than can lead to false detections. A 10 dB reduction in acoustic noise can be quite

useful because the desired signal is well above the detection threshold. The luxury of moderately strong signals in radio astronomy is all too rare.

In the past few years research into signal processing techniques for rejecting interference in radio astronomical receiving systems has become quite active. Initial results are encouraging, but they have also uncovered the complexities that must be understood and accounted for before applications are commonly available to observers. A few of the lines of pursuit are briefly outlined below.

5.1. Addition of information

As a rule, the information available in the radio telescope data stream is insufficient to distinguish interference from astronomical signals. Even the simplest data editing adds the bit of information that the interference is variable in intensity and the astronomical signal is not. If the interfering signal is too weak to detect with short integrations, this added information is insufficient. Any technique or *a priori* knowledge about RFI that enhances our ability to detect it will also increase our ability to reject it.

Similarly, any property of RFI that cleanly isolates it from cosmic signals offers the possibility of building a filter that rejects the RFI but passes the astronomical information. The following techniques use some form of RFI signal enhancement and signal isolation.

5.2. Possible techniques

Blanking On time scales much less than one second it becomes difficult to store data with sufficient time resolution to remove RFI with data editing. In that case data rejection in the time domain must be done in real time. Two examples where this is effective are on pulsed radar and Iridium satellite transmissions as mentioned in the “Spectrometer type and configuration” section. Radars typically transmit pulses of 2 to 200 microseconds in length with the pulses spaced by about 3 milliseconds. After about 500 microseconds most of the detectable echoes from nearby reflections have died away, and the time thereafter, until the next pulse, is available for astronomical measurements. The blanking generator must know the radar’s pulse spacing sequence so that it is only required to track slow drifts in the pulse generation clock. Blanking on individually detected pulses is ineffective because the weak pulses are below the detection threshold but add up to a significant feature in the integrated spectrum.

Blanking on time scales much less than one millisecond is problematic with current spectrometers because they require contiguous runs of data samples for intervals greater than the inverse of their frequency resolution. For example, 5 kHz resolution requires a 200 microsecond minimum data window. A spectrometer that works with non-contiguous data samples is conceivable, but none presently exist.

Radar and Iridium signals have frequency structure that is particularly bothersome to astronomical spectral lines, so these interfering signal must be suppressed well below the detection threshold. Some forms of broadband pulsed interference, such as lightning an ignition noise, does not have so much structure on frequency scales less than a few MHz so some residual interference may be tolerable. For this type of RFI pulse-by-pulse detection is necessary because of

the unpredictable pulse timing, but this may provide enough suppression to be useful, particularly if some antenna gain can be added to the detection receiver.

Cancellation Any signal in a radio receiver can be cancelled by finding a clean copy of the signal and adding it to the receiver signal path with the appropriate amplitude and reversed phase. One way to obtain a copy of an interfering signal is to receive it with a high gain antenna that has little or no response in the direction of the astronomical object being observed (Barnbaum & Bradley, 1998). The reference copy of the signal must have a relatively high signal-to-noise ratio to produce adequate suppression and to avoid adding noise to the system from the reference receiver channel. The reference signal must be processed by a fairly complex filter to replicated the delay of the interfering signal received in the main radio telescope receiver. In many cases the RFI signal in both the main and reference channels will arrive through several or even many propagation paths, each with its own delay and attenuation (Fisher, 2001). This complicated delay and attenuation must be duplicated in the reference channel filter. The relative phase and amplitude of the main and reference channel RFI signals will change with time, sometimes quite rapidly, so an adaptive feedback scheme is necessary to maintain cancellation.

Another method, called parametric cancellation, involves deriving a copy of the interfering signal from detailed knowledge of the transmitted signal. Then only a few, slowly varying parameters, such as intensity and delay, need to be solved for and tracked in real time. Ellingson, Bunton & Bell, (2000) showed this technique to be quite effective at suppressing the signal from one of the GLONASS satellites received in the sidelobes of the Australia Telescope Compact Array. The satellite signal had one clearly dominant propagation path which required a solution for delay, phase, Doppler shift, and amplitude roughly every tenth of a second. The GLONASS signal has a complex phase-switched signal structure, but the switching sequence is entirely predictable for 20-millisecond intervals. From the derived parameters a digitally synthesized version of the interfering signal was produced and added to the sampled telescope output to achieve RFI suppression below the detection level in ten seconds of integration.

Null steering A radio telescope is quite effective at suppressing RFI in the spatial domain by having 50 or 60 dB more gain in the direction of an astronomical source than in the direction of an interfering transmitter. Unfortunately, many interfering signals are considerably stronger than 60 dB above our detection threshold. In principle, one can create nulls in the telescope antenna pattern for selected interfering sources, either by adding the signal from one or more null-producing auxiliary antennas in the case of single-dish radio telescopes or by properly weighting the signal from an array of antennas that make up a radio telescope. In many ways the signal processing to produce spatial nulls is very much like the cancellation technique with a reference antenna described above. Many of the same signal-to-noise ratio and adaptive requirements apply. One extra factor that enters into steering nulls in an array is that the desired astronomical signal may also be suppressed unless extra restrictions on the null-steering algorithm are enforced.

The concept of null-steering is clearest when the interfering signal arrives from a well-defined direction. When significant multi-path propagation is present, as mentioned above, the required nulls are not simply in the spatial domain. They exist in the combined spatial-temporal domain. An interfering signal arriving from many directions is coherent with itself, but it may be cancelled only by the proper combinations of delay and amplitude from array elements or main telescope and auxiliary antennas. These combinations do not produce a directional null in space.

Post-correlation The equivalent to real-time signal cancellation using a reference antenna may be done with post-correlation data since the phase and amplitude of the main radio telescope and reference signals are preserved in the cross-correlation process. Briggs, Bell & Kesteven (2000) give the details of the RFI subtraction mathematics and show examples on data recorded at the Parkes radio telescope. One advantage of post-correlation interference subtraction is that the data may be processed off-line and iteratively optimized for the best cancellation parameters, if necessary. Aside from the correlation process itself, the data processing requirements, in numbers of arithmetic calculations per second, can be much less for post-correlation interference subtraction than for real-time adaptive cancelling since the correlation products may be integrated for as long as the interfering signal is stable in phase and amplitude, typically on the order of a second.

5.3. Where to from Here

There is no magic bullet for for RFI excision. For the foreseeable future a somewhat different approach may be required for each interfering signal type depending on how it is distributed in the time, frequency, and spatial domains as seen from the radio telescope. A technique that works well for pulsar observations may not be adequate for spectral line measurements. As we gain more experience with signal processing techniques for removing RFI some common tools will emerge. This is an interesting research topic, and I encourage anyone who is interested to join the effort.

A fairly comprehensive compilation of RFI mitigation work in radio astronomy may be found at the following web site:

<http://www.atnf.CSIRO.AU/SKA/intmit/>

References

- Barnbaum, C. & Bradley, R. F. 1998, A New Approach to Interference Excision in Radio Astronomy: Real-Time Adaptive Cancellation, *Astron. J.*, vol. 116, pp2598-2614
- Briggs, F. H., Bell, J. F. & Kesteven, M. J. 2000, RFI Subtraction from Time-Averaged Spectroscopy Using a Reference Horn and Cross Power Spectra, <http://www.atnf.csiro.au/~fbriggs/RFI/briggs.PS.gz>
- Ellingson, S. W., Bunton, J. D. & Bell, J. F. 2000, Cancellation of GLONASS Signals from Radio Astronomy Data, *Astronomical Telescope and Instrumenta-*

tion 2000 - Radio Telescopes, SPIE conference 4015, Munich, March 2000 (PDF version available at <http://esl.eng.ohio-state.edu/rfse/>)

Hagen, J. 1988, FAA Radar Tracker, Arecibo Observatory Electronics Department Manual No. 8806

Kesteven, M. 2000, PKS Interference-Excision Experiment, http://www.atnf.CSIRO.AU/intmit/srt/mbeam/kesteven/pks_test.html

Rabiner, L. W. & Gold, B. 1975, Theory and Application of Digital Signal Processing, Prentice Hall (Englewood Cliffs, N.J.)