

Linrad: *New Possibilities for the Communications Experimenter, Part 1*

Discussion opens with analog versus digital RF-input techniques and attendant performance considerations.

By Leif Åsbrink, SM5BSZ

In the old days, amateurs built their own equipment. Experiment was a natural part of the hobby. Inspired by what others did, one had to try to use the parts from one's own junk box to do something similar—or better.

With the commercial availability of modern amateur SSB transceivers, the need for experimentation declined. It is not easy to design and build equipment that can compete with commercial units. The best way to get a well-performing station has long been to buy a SSB transceiver. Only very few real enthusiasts build their own receivers and transmitters.

Today, the situation is changing. By use of simple equipment well suited for home building, radio signals can be moved into the digital world. Once a signal is available in digital form, a whole new field of experimentation is opened. A computer can do everything that we did before using analog electronics. Experimenting with different filter characteristics, AGC (automatic gain control) and AFC (automatic frequency control) can be done in soft-

ware at little cost—except the cost of the experimenter's time.

New possibilities in interference reduction treat different kinds of interference as separate signals, each one received with a digital receiver that optimizes the signal-to-noise ratio (SNR) for the particular interference. That brings with it a wide world of experiments with radio receivers, allowing reception of signals that cannot be received at all with a conventional SSB transceiver.

Linrad is a computer program that works under the *Linux* operating system on a standard PC (IBM-compatible, x86). This program receives an IF passband in digital form using simple read statements. Performance is determined by the hardware; the program just analyses the digital data. *Linrad* can be used to process the audio output from a conventional SSB radio or any other linear receiver with a bandwidth that can be handled by the sound card of the computer. *Linrad* can also use two audio channels (stereo) to process the *I/Q* pair produced by a quadrature mixer (direct-conversion radio). *Linrad* is designed for use with radio A/D converters in the future when boards and device drivers for *Linux* become available. The

Linrad program package is an ongoing development project.

In this first part of a series, I cover general receiver design philosophy. Part 2 will illustrate practical aspects of bringing radio signals into a computer via soundboards while subsequent articles will take up the specifics of *Linrad*, how to install it, what the special features are and how to use them to improve reception of weak signals.

Fundamental Receiver Operation

Fundamentally, there is no difference between analog and digital receivers. Both have the same basic problems with dynamic range and spurious responses. All the new digital methods for interference fighting have their equivalent analog counterparts. Let's begin with a general discussion of radio receivers intended to resolve common misunderstandings and explain how one can make sure a receiving system is properly optimized. The well-known problems of noise and dynamic range are central.

Linear Receivers

The ideal receiver for weak signals is the linear receiver. It uses linear processes only, and it may be imple-

mented in analog or digital hardware or by some combination thereof. The ideal receiver is completely quiet. That is, the output should be zero if a resistor at a temperature of zero kelvins (-273°C) were connected to the input. When something else is connected (an antenna), the ideal receiver selects a narrow part of the frequency spectrum, amplifies it and converts it to audio. That is all!

Linear processes may be split into several linear processes and applied after each other in any order. Because of the limitations of the available building blocks, one must use many linear processes to realize something that is near an ideal linear receiver.

Below is a list of the most important building blocks of a receiver for 144 MHz (see Fig 1). It is probably the most difficult band, in that very low noise figure is useful at the same time that very strong interference may be present, both in band and out of band. Notice that the goal is to come as close as possible to an ideal receiver so there are no compromises—just in case. If compromises are required, they should be made with full awareness of what the conflicting desires are, so no performance is thoughtlessly discarded.

For EME (moonbounce, reflecting signals off the moon) it is essential to have a receiver with performance close to that of an ideal receiver. When CW messages are received at a level where many repetitions are required, a small degradation in SNR will have large effects. Improving SNR by as little as 0.2 dB may be the difference between success and failure.

In “normal” communications, an improvement of 1 dB in SNR can barely be noticed. Degrading the noise floor by a single decibel makes it much easier to obtain good resistance against cross modulation and overload.

Noise Temperatures

Noise is best expressed as a noise temperature because noise temperatures are additive. A resistor that is kept at a temperature T will deliver the power $P = kT$ to another resistor with the same resistance for each hertz of bandwidth. (Where k is Boltzmann's

constant, 1.38×10^{-23} and T is in kelvins.) At room temperature, $kT = -174$ dBm/Hz. Two resistors having the same temperature are in thermal equilibrium and deliver the same power to each other. If one resistor is at 0 K, it will be heated by the power from the other resistor at temperature T according to the formula. The associated voltage is $V = \sqrt{kTBR}$, where B is the bandwidth in hertz and R is the source resistance in ohms. The equation is valid for normal temperatures and frequencies. For infinitely high bandwidths, V does not go to infinity because of quantum-mechanical effects.

The noise temperature of an amplifier is defined by a thought comparison with an ideal amplifier having exactly the same gain. It is the temperature that a resistor at the input of the ideal amplifier would have to produce the same noise power as the actual amplifier produces when connected to a resistor at 0 K.

In a system of amplifiers and other signal-processing hardware, the noise contribution from each stage can be expressed as a temperature component at the input. These temperature components are additive, and their sum is the system noise temperature.

The Input Amplifier

The input amplifier should not add significant amounts of noise to the signal received from the antenna, and it should not saturate from strong signals that may be present.

To obtain near-ideal noise performance from a receiver while maintaining good resistance to cross-modulation and overloading requires a good understanding of the problems involved. Look at ham.te.hik.se/~sm5bsz/pcdsp/preamp.htm, which mainly discusses the input circuitry and the tradeoffs between noise figure, selectivity and how they relate to the L/C ratio. The discussion is relevant to higher bands where too often, low impedance levels are chosen for the input filters. A 50- Ω transmission-line high-Q resonator may provide unnecessarily good filtering while it degrades the noise temperature too much. A higher impedance for the input filter

yields less noise and should be a good choice at 432 and 1296 MHz, where the low sky temperature makes system noise particularly important.

Once the compromise between noise performance and overload characteristics has been achieved, one should be sure the preamplifier is as close to ideal as possible. The famous Murphy's Law says something we should not forget: “Anything that can go wrong, will go wrong.” It is a good idea to place an overload detector at the output of all amplifiers that are followed by filters. The need for a filter directly after the preamplifier is discussed below.

The Second RF Amplifier

The gain of the preamplifier is usually insufficient to overcome the noise floor of the first conversion process, which is a frequency mixer in today's technology; but sometime in the future, it will probably be an A/D converter. Although it's possible to make broadband amplifiers with enough power output to preclude saturation by RF from the preamplifier, it is generally a good idea to insert some selectivity between the preamplifier and the second RF stage. Since this filter is mainly a precaution, there is no reason to make it complicated or to allow it to cause much attenuation in the passband.

The preamplifier is normally located very close to the antenna, while the rest of the receiver is placed indoors, so there is an attenuator in the form of a long cable between the preamplifier and the second RF stage. High preamplifier gain, low noise figure of the second RF amplifier and modest losses are required for nearly ideal performance. Table 1 shows data for 144 MHz. The noise figure of the second RF amplifier includes all losses between the preamplifier and the second amplifier. Attenuators inserted before an amplifier degrade the noise figure with their attenuation.

Table 1 shows degradation caused by the second RF amplifier, which is treated including the NF degradation caused by cable and filter losses. T is the temperature associated with the NF at the second amplifier input. $T(\text{ant})$ is the same noise contribution




Fig 1—A block diagram of a typical receiver front end to feed a sound card for DSP.

when referenced to the input of the first preamplifier.

The noise temperature caused by antenna and preamplifier are assumed as shown in Table 2, with two alternatives. To obtain good dynamic range farther down the signal path, it will be necessary to allow some contributions to system noise from other amplifier stages. To keep the total excess noise low, the contributions must be very small. An inspection of Table 1 shows that even on 144 MHz, where the antenna temperature is not very low, the gain of the preamplifier must be about 20 dB.

To allow 3-dB of cable/filter loss and a second RF amplifier noise figure (NF) of 2 dB, preamplifier gain must be 25 dB. A neutralized GASFET with power-matched output has a gain on the order of 30 dB, and it allows 8 dB of combined filter and cable attenuation if the second amplifier has a NF of 2 dB.

Such a receiver front end does not have an optimum third-order intercept point (IP3) for in-band signals, but performance is usually good enough for practical purposes. A two-tone test of such an amplifier with an MGF1801 shows a mediocre IP3 of 0 dBm at the input. Signals up to about -30 dBm can be allowed without serious problems, since two -30-dBm signals give third-order IM spurs (IMD3) corresponding to -90 dBm at the input. For in-band signals, IM3 at such levels is not likely to be a problem. The noise floor is at -174 dBm/Hz, and so as not to destroy weak-signal operation, the interfering stations must have their noise sidebands below -145 dBc/Hz. Amateur Radio equipment is not quite that good as far as I know, so desensitization dynamic range (DDR) will be the limiting factor.

Should it turn out to be desirable, it is not difficult to improve the in-band IMD3 performance of the preamplifier by using so-called noiseless feedback stages and a second preamplifier using a device running at higher power, also with feedback. Table 3 shows the characteristics of a nearly ideal receiver based on the output-power matched, neutralized, MGF1801.

From the data in Table 3, we can get the noise temperature at the input of the second RF amplifier, referenced to the antenna input: $T_{rx} = 34231 / (0.315 \times 500) = 217.3$ K. The contribution from the second stage is 2.3 K.

The amount of gain required in the second RF amplifier will depend strongly on the noise figure of the next stage. The gain and output intercept point required for the second RF amplifier is listed in Table 4 for

different assumptions of the noise figure for the next stage, usually a Schottky-diode mixer.

Amplifiers having noise figures of 2 dB and power outputs up to 2 W are not difficult to design. Gain and saturation power level depends on the noise figure and maximum input power for the first frequency-conversion stage.

The RF Band-Pass Filter

The first conversion, be it an A/D converter or a diode mixer, has spurious responses. A filter with adequate suppression for signals on the spurious frequencies must be inserted in the signal path before the first conversion stage.

If necessary, the RF band-pass filter can be made very narrow with rela-

tively high attenuation in the pass-band. Any attenuation caused by the band-pass filter adds to the noise figure of the third stage, which leads to a bigger transistor in the second RF stage, as shown in Table 3.

When using RF A/D converters with today's technology, the sampling speed is typically 50 to 100 MHz. All frequencies between 0 and 200 MHz fall between zero and half the sampling frequency. The RF band-pass filter must allow only one set of alias frequencies, in a bandwidth of half the sampling frequency, to reach the A/D converter.

When using a local oscillator and a frequency mixer for the first conversion, the RF filter must suppress not only the mirror frequency but also frequencies that would mix with the LO overtones and give an output at the

Table 1—Degradation from a Second RF Amplifier Including Cable and Filter Losses

T is the temperature associated with the *NF* at the second-amplifier input. *T*(ant) is the same noise contribution when referenced to the input of the first preamplifier.

Preamp Gain (dB)	Second Amplifier + Loss			S/N Loss	
	NF (dB)	<i>T</i> (K)	<i>T</i> (ant) (K)	At 215 K (dB)	At 272 K (dB)
15	1	75	2.4	0.04	0.03
15	2	170	5.4	0.11	0.09
15	3	290	9.2	0.18	0.13
15	4	439	14	0.27	0.22
15	5	627	20	0.39	0.31
15	6	870	28	0.53	0.41
20	1	75	0.8	0.01	0.01
20	2	170	1.7	0.03	0.03
20	3	290	2.9	0.06	0.05
20	4	439	4.4	0.09	0.07
20	5	627	6.3	0.12	0.10
20	6	870	8.7	0.17	0.14
20	10	2610	26	0.50	0.40
25	1	75	0.2	0.00	0.00
25	2	170	0.5	0.01	0.01
25	3	290	0.9	0.02	0.02
25	4	439	1.4	0.03	0.02
25	5	627	2.0	0.04	0.03
25	6	870	2.8	0.06	0.05
25	10	2610	8.2	0.16	0.13
30	10	2610	2.6	0.05	0.04

Table 2—Noise Temperature Assumptions

Source	Noise 1	Noise 2
Sky	167 K	167 K
Sidelobes	15 K	40 K
Antenna losses	5 K	5 K
Cable + relay	13 K (0.2 dB loss)	30 K (0.4 dB loss)
Preamplifier	15 K (0.22 dB NF)	30 K (0.4 dB NF)
Total	215 K	272 K

IF. A narrow RF filter will also suppress false responses caused by spurs that may be present in the LO signal. There are more reasons for a narrow RF filter below.

The First Conversion: A/D Converter or Diode Mixer

Someday, an A/D converter will be the first mixer in receivers up to many hundreds of megahertz. Today's technology allows 65-MHz sampling at 14-bit resolution using, for example, the AD6644 from Analog Devices. Such a chip offers typically 74 dB of SNR at a bandwidth of 32 MHz, corresponding to 149 dBc/Hz. The peak-to-peak amplitude such a device needs for full range is 2.2 V, corresponding to about 11 dBm into a 50-Ω load. The noise floor is at -138 dBm/Hz, which corresponds to a noise figure of 37 dB, all referenced to a 50-Ω load. The input impedance of the AD6644 is 1 kΩ, so the power actually consumed by the chip is -2 dBm for full scale. If a lossless impedance transformer were used at the AD6644 input, rather than a 50-Ω termination resistor, the noise floor would be at -151dBc/Hz, with an associated noise figure of about 23 dB.

The AD6644 receives noise from nearly 10 times more bandwidth than represented by the digital bandwidth. It is not possible to lower the noise figure by more than about 6 dB without using a selective amplifier or degrading the dynamic range. Allowing 3 dB of loss for the RF filter, Table 4 (extrapolated) shows that the AD6644 will need 39.6 dB of gain in the second RF stage if the AD6644 is terminated in a 50-Ω resistor. Saturation of the AD6644 will occur at 14 dBm, 38.6 dB below the level where the MGF1801 saturates.

High-level Schottky-diode mixers (level-23 from Mini-Circuits) have IP3s around 30 dBm and 1-dB compression points of around 15 dBm. Consider a 10-dB gain second RF amplifier using some transistor that can deliver 200 mW (23 dBm) through the RF filter with an impedance-matching network required by the mixer. This will make a level-23 mixer saturate just before the preamplifier, with an IP3 around -2 dBm at the input. Since mixers attenuate by about 8 dB and the noise figure at the mixer input has to be maximum 10.4 dB, the amplifier after the mixer must be very quiet lest it degrade the system noise figure. Compensating a poor IF-amplifier noise figure with more gain in the second RF amplifier will degrade system intercept point.

It follows from the discussion above that the MGF1801 power matched and neutralized amplifier is good

enough when it comes to power-handling capabilities. A 10-dB improvement is not difficult by noiseless feedback, but to utilize the improvement one must design very special low-noise, high-level mixer/IF combinations. In cases where signals well outside the desired passband cause the preamplifier to go nonlinear, noiseless feedback and additional filters and amplifiers may be needed.

The A/D converter has superb linearity. IP3 is only about 8 dB lower than a well designed level-23 diode mixer. A diode mixer can be driven into the nonlinear region by a single very strong interfering signal without any serious degradation of a desired weak signal, while the D/A converter produces useless data when saturated. Saturation occurs about 36 dB earlier in an AD6644, so if there is only one interfering signal, the diode mixer allows about 30 dB more interference signal. It is quite clear already that today's A/D

technology is extremely attractive.

The First Local Oscillator

If the first conversion is an A/D converter, the first LO is the sampling clock, which has a fixed frequency. The first LO will be at a fixed frequency also were a diode mixer used in conjunction with a broadband IF. Typically, an LO frequency of 116 MHz is used to convert 144 MHz to 28 MHz.

Fixed-frequency oscillators using crystals can be made with very little phase noise. It may seem very simple to use a 12.88888-MHz crystal and two frequency triplers to produce 116 MHz. Good filtering is required, though, because the thirteenth overtone of 12.88888 MHz is at 167.55 MHz, which may cause a spurious response at 139.55 MHz that may not be suppressed much by the RF filters. Good filters, with two L-C circuits, to make the output of all frequency-multiplier stages pure will prevent this problem.

Table 3—Typical Output-Power Matched MGF1801 Preamplifier (Negligible System Noise Degradation)

Preamplifier

Antenna temperature = 200 K
 Preamplifier NF = 0.2 dB = 15 K
 Preamplifier gain = 27 dB = 500 times in power
 Noise temperature at output of preamplifier = (200 + 15) × 500 = 107500 K
 Input intercept point = 0 dBm
 Saturated power output = 18 dBm

Cable/Filter

Losses = 5 dB (gain = 0.315 times in power)
 Output noise temperature = 0.315 × 107500 + (1 - 0.315) × 290 = 34061 K

Second RF Amplifier

NF = 2 dB = 170 K
 Noise temperature at input = 34061 + 170 = 34231 K
 Preamplifier intercept point at input = 0 + 27 - 5 = 22 dBm
 Saturated power input = 18 - 5 = 13 dBm

Table 4—RF Amplifier 2 Requirements

Gain, Output Power and Output IP3 required of RF Amplifier 2 for different Noise Figures of the Third Stage to make the Third Stage Contribute 2 K at the Antenna Input. This table is based on the data of Table 3.

Saturated Minimum Output				
NF (dB)	Temp (K)	Gain (dB)	Power Output (dBm)	IP3 (dBm)
6	865	4.4	17.4	26.4
9	2013	8.1	21.1	30.1
12	4307	11.4	24.4	33.4
15	8880	14.5	27.5	36.5
18	18009	17.6	30.6	39.6
21	36221	20.6	33.6	42.6

When the first IF is routed to a narrow filter at typically 9 MHz or 10.7 MHz, the first LO must produce a variable frequency. It is very difficult to make good oscillators with variable frequency. The first LO is usually the limiting factor for receiver dynamic range in well-designed receivers having a first IF of narrow bandwidth. Many articles in the amateur literature describe the design of low-noise frequency synthesizers for LO use.

It is not easy to select a good frequency for the first local oscillator. There are many problems. For example, using 116 MHz to convert 144 MHz to 28 MHz has the following problem: A signal at 145.3 MHz will produce its main signal from the IF port at 29.3 MHz. If the signal were very strong, the third overtone of the IF signal at 87.9 MHz would be present inside the mixer. There, it would be mixed with 116 MHz to produce a false signal at 28.1 MHz that, in turn, would cause interference for signals at 144.1 MHz. There will always be combinations of overtones of the IF signal and the LO or its overtones that would fall within the IF passband, causing spurs at some frequencies. To avoid such problems, it is a good idea to avoid LO/IF combinations that give spurs like this of low order, and it is a good idea not to make the RF passband wider than necessary.

Broadband IF Filters and Amplifiers

The problem of amplifying and filtering the signal present at the output of the first mixer is identical to the problem of designing the RF amplifier and RF filter section. The noise figure need not be as low: The only reason to have a very low noise figure is to get good dynamic range in earlier stages, because their gain can be made lower.

If the IF amplifier had a noise figure of 0.6 dB (= 43 K), the stage limiting the in-band dynamic range will be allowed to contribute with 127 K for the IF noise figure to become 2 dB. If the IF amplifier has a noise figure of 1.6 dB (= 130 K) the gain of the IF amplifier must be increased by 5 dB to keep the IF noise figure at 2 dB. That would be a bad idea: It is better to increase RF gain and accept a worse IF noise figure.

The purpose of the broadband IF filter is to suppress the mirror frequency and the spurious responses of the next frequency conversion. The first broadband IF filter may also be used to shape the pulse response of the receiver, to allow an efficient noise blanker. A wide bandwidth with roughly Gaussian frequency response makes interference pulses very short

and allows an efficient noise blanker.

Conversion to Baseband

The conversion to baseband is normally incomplete in analog receivers. The baseband signal is a complex signal that has two components, in-phase (*I*) and quadrature (*Q*). The two signals, *I* and *Q*, contain the same frequencies and their phase relation contains information about whether the signal at baseband is above or below the frequency of the last LO (the BFO). In analog filter-based SSB receivers, one uses a narrow filter—the last IF filter—to ensure that no signal is present on one side of the last LO (undesired sideband). Thus, one is sure that any signal present at baseband must be from the desired sideband. Consequently, there is no need to produce both *I* and *Q* because their phase relationship will give no new information.

An analog receiver for AM may convert the IF signal to a baseband signal with both *I* and *Q*. If the frequency of the LO were very close to the carrier of the AM signal, the *Q* signal could be used to control the last LO frequency through a low-pass filter. This way, the LO becomes phase locked to the carrier and the modulation is in the *I* signal only. The noise in the *Q* channel will not contribute and some improvement of SNR is possible.

Digital-signal processing is easier and more efficient at baseband with complex signals. There are several different ways to go from RF to the digital baseband *I-Q* pair. When the RF signal is fed to a AD6644 sampling at 65 MHz, a 144-MHz signal will be at 14 MHz in the digital data stream. A general-purpose DSP or a modern PC is not fast enough to process data at 65 MHz. There is a chip, the AD6620 (Analog Devices), that mixes frequencies to baseband by multiplying the input samples by a sine/cosine function at a frequency specified by the user.

The AD6620 is normally used to sample the RF signal directly so the input data stream is real data. To listen to 144 MHz for weak signals, one would make the internal digital oscillator of the chip operate at 144.15 MHz, typically. The 14.000-MHz signal corresponding to 144.00 MHz is then converted to two signals at 150 kHz, with a phase shift of 90° between them. A signal at 14.300 MHz, corresponding to 144.3 MHz, would also be converted to two signals at 150 kHz; but with the opposite phase shift. Because of the precise phase and amplitude relations possible in digital circuits, these two signals can be completely separated in later processing stages.

The AD6620 contains decimating

filters that gradually bring the sampling speed down. Unfortunately, these filters do not allow very steep cutoff so the spur-free bandwidth becomes very limited in relation to the re-sampled data speed. There are better data decimation chips, for example the GC4016 from Graychip, which will allow extremely steep filters and an output data stream that is not over-sampled but still spur-free.

I have no practical experience with AD6644 and AD6620 or GC4016. These chips would allow more bandwidth and better dynamic range compared to most sound-card-based solutions.

When using an audio board to convert from analog to digital, there are two ways to go. One is to filter a well-defined passband by means of an IF filter with steep skirts that will allow very good suppression of the image frequency on the other side of the LO. In this case, a single mixer and one A/D converter channel is required for each RF channel. With a standard audio board, sampling at 44.1 kHz, one can receive two independent signals this way at bandwidths up to about 20 kHz. Unless great care is taken, overtones in the audio range (IM2) cause strong spurs that can easily be avoided by allowing only a 10-kHz bandwidth, in which case the audio range is placed typically from 10 to 20 kHz.

When the A/D converter is sampling real data (the filter method), the conversion to complex data is made in the computer. The other way is to produce the baseband complex pair (*I* and *Q*) in analog hardware (direct-conversion radio), in which case two audio channels are needed for each RF signal.

Creating the baseband signal in analog hardware saves some computer time and gives more bandwidth—twice as much—because of the use of two audio channels. It is advantageous to eliminate sophisticated IF filters, but direct conversion has the problems of audio overtones and requires great care for extremely good linearity of mixers and audio amplifiers. The local oscillator must be extremely stable, because modulation on it will be detected (mixed with the LO carrier) to produce interference in the audio range.

Summary

This article has concerned itself with the broad issues involved in receiver design. In my next segment, I shall discuss details in design of the hardware required to move radio signals into the computer by use of sound-cards. □□