Linrad *with High-Performance Hardware*

Together with the WSE RX converters, Linrad *is a software-defined receiver that should exceed any other receiver in dynamic-range performance.*

By Leif Åsbrink, SM5BSZ

Linrad has evolved from earlier systems that I have worked with since about 10 years. My main interest in Amateur Radio has always been the technology for weak-signal communication. In 1993 I erected a cross-Yagi array, 4×14 elements and started working EME on 144 MHz. Being able to eliminate Faraday rotation turned out to be very efficient but not so easy on extremely weak signals. I needed a computer to assist.

The first version of what is now *Linrad* was implemented on a TMS320C25 system. This system could display a 3 kHz wide window on an oscilloscope as the summed power spectrum from both the polarizations. With an averaging time of a few seconds and about 10 Hz resolution, signals could be seen before they were possible to copy. In the 1995 ARRL EME contest, the TMS320 system was capable of locking to a signal, filtering it through 17-Hz band-pass filters and combining the two signals from the two orthogonal antennas automatically to produce the optimum fit to the polarization of the incoming wave. The signal produced by a receiver automatically keeping the filter centered and the polarization aligned was then sent to my headphones. I scored number two, after W5UN, in the single-operator class that year and was very happy with this system even though it was completely inflexible, with all code in an EPROM that had to be produced on another system. With a 12-bit A/D converter, this system had a poor dynamic range, so it was completely saturated when a local station entered the passband. Having to keep SM5FRH, SM5DCX and a few others outside a "3-kHz window" most of the time was of course a limitation for this system, but not a serious one.

In 1997, I started to move the algorithms into the PC and had a working system in 2×20 kHz bandwidth about one year later. This system was under MSDOS and the increased bandwidth made operation much more exciting. The SoundBlaster 16-bit A/D converters allowed a much better dynamic range; this system was only saturated for about one hour at moonset when SM5DCX had his main lobe straight into my back lobe. (My antenna was an extreme "maximum-gain design" with a large back lobe.)

Having software running in a real computer, with all the flexibility coming along with that, made it possible to analyze the EME signals better. As it turned out, "144-MHz EME signals" are only about 0.25 Hz widened by multi-path propagation, so introducing coherent processing was an obvious thing to do. By the time I found that coherent averaging is possible and started to include routines for that, I found that I had to restart the entire project because the code was

Jaders Prastgard 3265
63505 Eskilstuna, Sweden
leif.asbrink@mbox300.swipnet.se

becoming too messy and my homemade drive routines were becoming obsolete. This was in 2001, when I made a new start under *Linux*. This time I had a much better idea about what I wanted the program to do. First of all, I wanted flexibility and hardware independence. By now, autumn 2003, *Linrad* contains everything that was ever included in the older systems but not much more. Many more things are planned for the future, but at present my focus is on high-performance hardware to use with *Linrad*. In recent years, EME activity has spread out over a much wider frequency range—20 kHz is no longer enough. Terrestrial communication also calls for more bandwidth. More bandwidth calls for more dynamic range. It is not possible to keep strong local stations outside the passband, so the D/A converter must have the dynamic range required to handle very strong signals. This article focuses on the system I am currently putting together for 144 MHz EME. *Linrad* and the hardware is in no way limited to this usage, it just happens to be at the focus of my own personal interest and technically it is a very demanding mode of operation.

The WSE Converters

To go from 144 MHz or other amateur bands to a digital data stream, I am using several cascaded converters. This may seem very complicated, but in a way it is not. A complicated problem is split into several less complicated problems. Each converter is doing just one simple task. It is well matched to 50 Ω at both input and output. Each unit can be evaluated separately, and it is not difficult to find out what the limiting factors are.

The WSE RX converters are designed for low noise and low gain. They are open designs, described in detail at **www.antennspecialisten.se/~sm5bsz/linuxdsp/optrx.htm**. The entire system is kind of a brute-force solution to the problem of receiver dynamic-range limitations. Each converter uses about 18 W, mainly for the class-A buffer amplifiers, so the boxes must be rather big to provide a low temperature without forced air cooling. The design uses through-hole-mounted components only and is experimenter-friendly. Assembled and tested units are available from **www.antennspecialisten.se**.

For 144 MHz, the four converters listed in Table 1 are used after one another. The system does not have any VFO, only low noise crystal oscillators, so dynamic range is not limited by reciprocal mixing. Each converter has two channels with a common local oscillator. Each channel has an RF input amplifier, an RF filter, a mixer and an IF output amplifier. The amplifiers have low gain, typically 8 to 10 dB, and they use noiseless feedback. In order to get some isolation, an attenuator follows each amplifier. Noiseless feedback transfers the output load impedance of an amplifier to the input. With 3 or 4 dB in each attenuator and 1 or 2 dB attenuation in the filter, there is enough gain to overcome the conversion loss of the mixer and provide between 0 and 4 dB gain, different for different units. There will also be at least one RXHF unit built in a similar fashion to convert from the HF bands to 70 MHz. The RXHFA will probably work for 1.8, 3.5, 7, 10 and 14 MHz. At present, the RX144 unit is in a late prototype stage, while RX70, RX10700 and RX2500 are available.

The WSE receive system is about 20 dB better than a conventional transceiver. This is a bigger difference than one can really use on the air, because the lack of spectral purity of the interfering station(s) will be the limiting factor. It will be possible to produce a transmitter the same way and get a similar transmit performance. There are several other ways to make an ultra pure transmitter, at least in CW mode. *Linrad* with the WSE converters is an excellent spectrum analyzer to use when building high-performance transmitters or when modifying standard transceivers for better transmit performance.

It should be obvious that far simpler solutions than the WSE RX converters will be adequate at most locations. I have made the WSE products for my own use. I will make a limited number of units available, and if demand is sufficient, there will be a continued supply from Antennspecialisten. Software-defined radios have different characteristics than conventional receivers. There will probably be SDRs available with seemingly good performance data that do not perform well when compared to a "good old analog" radio. The reason would not be the digital technology as such; the way dynamic range is specified may be misleading. IP3 is one of the commonly used figures of merit for receivers and it is discussed in some detail below. An analog radio will typically work fine with instantaneous voltages up to about 20 dB below IP3 while a digital one may become overloaded 40 dB below IP3 or wherever the A/D converter saturates. Together with the WSE RX converters, *Linrad* is a software-defined receiver that should outperform any other receiver when it comes to dynamic range. There is no limitation in the digital technology as such. Problems may arise when an A/D converter is fed with a large bandwidth because the instantaneous voltage caused by the summed amplitude of many signals may occasionally exceed the A/D converter range, and the conventional way of measuring receiver dynamic range might not show the limitations.

The A/D and D/A converters

The output from the RX2500 is four audio signals with a bandwidth of nearly 48 kHz each. To sample them, an A/D-converter with four channels and 96 kHz sampling speed is required. The second article of this series[1] gives some information about the RX2500 unit and the modified Delta44 sound card that I use to sample the four audio channels.

Better sound cards are available now, and replacing the Delta44 would improve the performance of the entire system. Someday, I hope someone else will determine what cards are best. Once the proper drive routines are installed, *Linrad* should work automatically.

The Delta44 uses the same speed for input and output. There is no reason at all to produce the output at a sampling rate of 96 kHz. *Linrad* is not written for that and the current code would be extremely inefficient. I

[1]Notes appear on page 31.

Table 1

With these four converters and a Delta44 sound-card, a 90 kHz wide passband at 144 MHz is converted to a digital data stream inside *Linrad*. The center frequency can be selected anywhere between 143.975 MHz and 145.975 MHz in steps of 25 kHz.

Name	*Input (MHz)*	*Output (MHz)*	*Crystals*	*Crystal Separation (kHz)*
RX144	144	70	4	500
RX70	70	10.7	5	100
RX10700	10.7	2.5	4	25
RX2500	2.5 baseband		1	—

use a standard audio card for the output at a sampling speed of 5 kHz for CW modes and 8 kHz for SSB. Internally in *Linrad*, the sampling rate at the output of the final filter is not higher than required for the bandwidth, an EME signal that has passed a 20 Hz filter is typically sampled at 46.875 Hz (96 kHz divided by 2048). In the final processing step, the signal is resampled by a fractional number to fit the output speed of the D/A converter. The signal is also frequency shifted by the BFO setting.

The output is kept synchronized with the input by gradually changing the fractional resampling rate. Since separate crystal oscillators generate the input and output sampling rates, the resampling rate will change with time. The total amount of sampled data points waiting in the various processing stages should correspond to a constant time. By monitoring the total processing delay it is possible to detect the need for a resampling rate change.

Linrad Setup: FFT Versions, Sizes and Windows

Assuming a cross-Yagi array and preamplifiers with adequate gain connected to the RX144, a system optimized for 144 MHz EME will need a waterfall bin bandwidth somewhere between 1 and 10 Hz. A good noise blanker is essential in most locations so the second FFT must be enabled. Running two channels at a processing bandwidth of 96 kHz requires a Pentium III or better, so version 5 should be selected for the first FFT. This is the fastest floating-point implementation, which uses the SIMD instructions (single instruction multiple data) to compute the transforms of both channels simultaneously.

The processing delay through *Linrad* is long, up to 10 seconds, for optimum readability of weak EME signals. This has nothing to do with processor speed, it is a consequence of the character of the EME path and the optimum parameters for the AFC. This means that there is no reason to select a small size for the first FFT to minimize processing delay. Adding 0.2 seconds by making the first FFT band-

Table 2

A 144 MHz preamplifier will lower the system noise figure. Assuming a noise figure of 0.2 dB for the preamplifier and 11 dB at the RX144 input, total system noise figure and dynamic range depend on the preamplifier gain as given by this table. The antenna temperature, Tsky is assumed to be 200 K and S/N loss is relative to an ideal (noise-free) receiver.

Gain (dB)	*NF (dB)*	*Temp (K)*	*S/N Loss (dB)*	*Dynamic-Range Loss (dB)*
0	11.0	3561	12.51	0
3	8.36	1898	9.77	0.26
6	5.97	1057	7.22	0.71
9	3.99	637	5.03	1.52
12	2.50	426	3.28	2.77
15	1.50	320	2.04	4.53
18	0.90	267	1.25	6.74
21	0.57	241	0.80	9.29
24	0.39	227	0.55	12.04
27	0.29	220	0.41	14.90
30	0.25	217	0.36	17.85
33	0.22	215	0.31	20.80

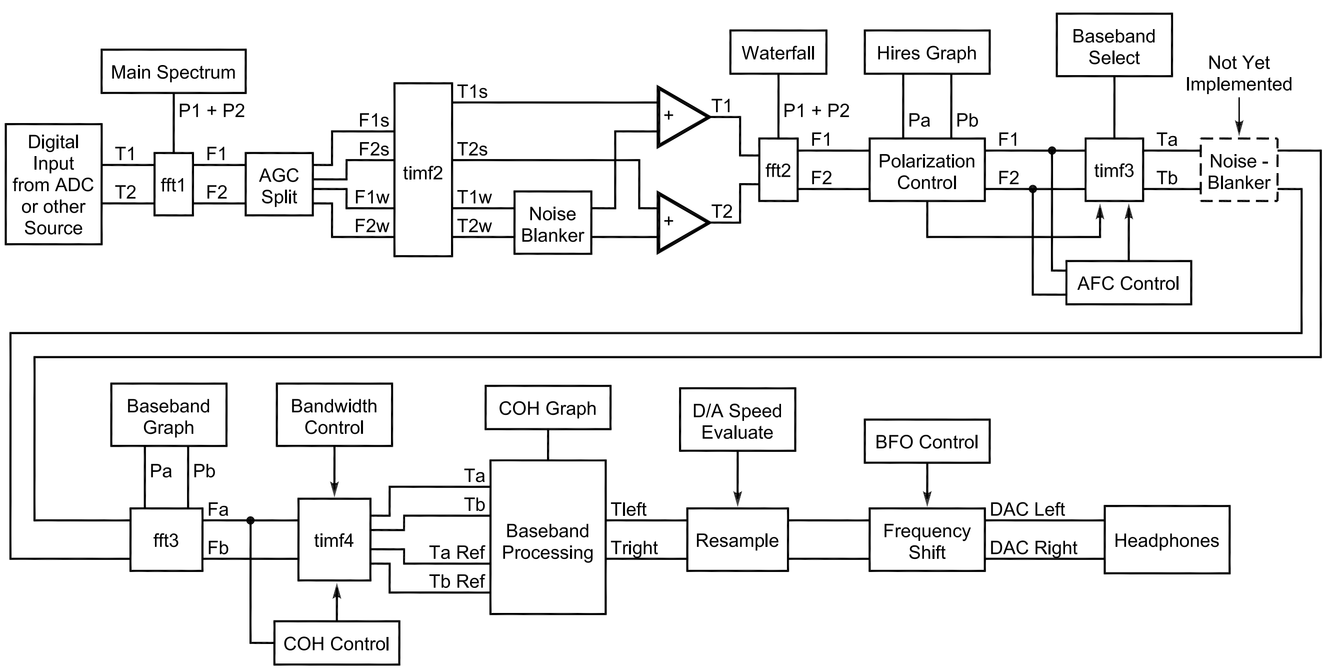

Fig 1—The block diagram of *Linrad* with two receive channels and the second FFT. T1 and T2 are signals in the time domain from two antennas 1 and 2. F1 and F2 are the corresponding signals in the frequency domain. Ta and Tb are linear combinations of T1 and T2 that make the desired signal zero in Tb and consequently maximizes the desired signal in Ta. Ta-ref is a time function constructed from a much narrower bandwidth than Ta. For Morse coded signals, it will be the CW carrier that is useful for coherent processing.

width 10 Hz is no significant disadvantage. Keeping a modest ratio between the sizes of the second and the first FFT makes it easier to ensure that very strong signals will not saturate the second FFT even if they are stable enough to put nearly all their energy in one single frequency bin.

Typical parameters would be a first FFT bandwidth of 30 Hz and second FFT bandwidth eight times narrower. The parameter is in powers of two, so it should be three that is 2^3. With large transforms, a window of $\sin^2$ is sufficient for the first FFT and for the second FFT the sine function itself ($N = 1$) is perfectly adequate. With *Linrad*-01.01 and later, these parameters will give the size 8192 for the first FFT and 65536 for the second with bandwidths of 23 and 2 Hz, respectively. The transform sizes come in powers of two, so you never get exactly what you ask for. The resampling spurs surrounding a very strong signal disappear into the noise about 2 kHz away from and 145 dB/Hz below a near saturating carrier with these parameters. By setting the first FFT window to $\sin^4$, it is possible to eliminate these spurs completely. They then disappear into the phase noise of the 2.5 MHz test oscillator 140 dB/Hz below the carrier at a frequency separation of 200 Hz. That would be a waste of CPU power because no interference source could be expected to have a spectral purity anywhere near –145 dB/Hz as close as 2 kHz.

Linrad Setup: FFT Signal Levels

First of all, the gain of the analog hardware should be set for the desired compromise between dynamic range and system noise figure. With the WSE converters, "setting the gain" is simply setting the gain of the 144 MHz preamplifier. With a system noise figure of 11 dB at the RX144 input and with a preamplifier noise figure of 0.2 dB, the in-band dynamic-range loss, system noise figure and preamplifier gain relate as illustrated in Table 2 for an antenna temperature of 200 K.

Table 2 shows the usual thing. One wants the preamplifier to lift the noise floor by something between 10 and 20 dB for a compromise between dynamic-range loss and noise figure. Dynamic-range loss is the amount by which the noise floor is lifted when the preamplifier is connected. The WSE converters, using only crystal oscillators, are not much affected by reciprocal mixing, so the dynamic range is the distance from the noise floor to a fixed power level where something becomes nonlinear.

As can be seen from Table 2, really low noise figures require high gain and will degrade the dynamic range by nearly 20 dB. In cases where dynamic range is the limitation, a preamplifier gain of 12 dB only will provide a noise figure of 2.5 dB, which will degrade an EME signal by 3.3 dB for an antenna pointing towards cold sky. For terrestrial modes, an antenna temperature of 1000 K is often assumed, in such cases even less gain could be considered.

The block diagram of *Linrad* is reproduced here as Fig 1. The major processing blocks are fft1, timf2 and fft2. These blocks compute forward, reverse and again forward FFTs at the full sampling rate. The design of a digital receiver is no different from the design of an analog radio. Each processing block has a saturation level and a noise floor. In the digital world one can make the dynamic range extremely large by use of many bits for each data point, but that has a penalty in CPU load. The 16-bit multimedia instructions run three times faster than floating point and therefore 16 bit data is used for timf2 and fft2. This leads to several complications, but computers were not fast enough when I wrote the code. There are several compromises in the *Linrad* architecture that may be removed in the future when CPU speed is no longer a limitation. The 16-bit processing blocks do give a small contribution to the system noise floor and they may limit the performance of the smart noise blanker. Going from 16 to 32 bit data words could improve dynamic range by a few tenths of a decibel, but spare CPU capacity may be used for many interesting things, so I have no plans for a change in the near future.

Fft1 must use 32-bit data to handle the full dynamic range. The output of fft1 is split into two blocks and an AGC makes sure no signal is strong enough to saturate when converted to 16 bits. The maximum level of the output from the AGC depends on several factors. The attenuation to use at frequencies where strong signals are present is calculated from power spectra. Three different power spectra are used for this purpose: A fast and a slow fft1 spectrum and a fft2 spectrum. The fast fft1 spectrum is intended to prevent overflows when a very strong signal starts suddenly. The averaged spectra are needed to find weaker signals that may be strong enough to degrade the noise blanker but do not have S/N enough to be found in a single fft1 power spectrum. A relatively strong signal may be hidden in the pulse noise that the blanker will remove and reasonably good statistics are required to find it. The fft2 spectrum does not have this problem, but there are some stability problems in using it because of the way *Linrad* is designed. The interference that will not be removed from a frequency on which there has been a strong signal can be interpreted as a strong signal if the blanker controls are used carelessly. For details about this phenomenon, look at **antennspecialisten.se/~sm5bsz/linuxdsp/blanker/leonids.htm**. When the fft1 bandwidth is as narrow as 23 Hz, it is a good idea to use unaveraged power spectra for the fast fft1 spectrum. Use the little box in the lower right corner of the main spectrum to set the number of spectra for the first average. Set it to one, the default value is five. Using unaveraged power spectra will cost some CPU time. It is necessary to do the averaging in two steps when the fft1 bandwidth is very large, but with the parameters given here the increased CPU load should not be a problem. Since the transform size is eight times bigger for fft2 than for fft1, strong signals that occupy one bin in fft1 only must be limited to eight times less power than the saturation limit. In the worst case, when all the energy comes in a single frequency bin in fft2 too, the energy is collected over an eight-times longer period.

Exactly as for analog processing blocks, it is essential that the noise floor is placed correctly for the digital processing blocks. 16 bits is marginal for the dynamic range needed. The WSE converters add 0.5 to 1 dB, each, to the system noise floor. Timf2 and fft2 add a few tenths of a dB each, as will be discussed below. The weakest link in the signal processing chain is the Delta 4 A/D converter. Despite the modification that lowers the noise floor by typically 3 dB, the Delta 44 produces about 40% of the system noise floor at the RX144 input. The system noise figure of 11 dB at the RX144 input is due to the summed effect of all the noise sources. The noise figure of the RX144 itself is about 6 dB.

When you start *Linrad* for the very first time, you are prompted to setup routines. Select the appropriate parameters for your sound card and enter a receive mode. You are again prompted for parameters, select the default ones or something else that seems appropriate. After the last parameter screen you get to the normal processing routine. Press "A" to make *Linrad* show amplitude information. The lower left corner of the screen will look like Fig 2. None of the values should become zero under normal operation. The numbers hold the minimum value and they may become zero due to the transient that may occur in case the A/D converter is stopped and restarted. They may also become zero at extreme events like changing the local oscillator frequencies while a very strong signal is present. Clear the minimum value by pressing "Z." If any of the numbers tends to become zero often, some signal level is too high.

These parameters are digital volume controls:

"First FFT amplitude" is fft1 input.
"First backward FFT *att. N*" is timf2 output.
"Second forward FFT *att. N*" is fft2 output.

These volume controls affect the signal levels inside the major processing blocks. The 16-bit processing blocks timf2 and fft2 are the critical ones. The volume controls should be set for the timf2 and fft2 noise contributions to become negligible. The dominating contribution for timf2 is the rounding error in going from floating point to integers, about 0.3 bits RMS. Since the rounding errors at the timf2 input are made in the frequency domain, lowering the input volume control for the quantization noise to grow to a substantial fraction of the noise floor does not lead to a S/N decrease. It works the other way around. The signal becomes enhanced! Not very surprising at second thought because when all the frequency bins containing only noise have amplitudes below one bit, the noise disappears completely. The signal will not disappear if its amplitude is above one bit. This is an artifact. When back transformed, such a signal is distorted and completely useless if it is near the noise floor. To really verify the S/N loss caused by rounding errors, the signal must be well below the noise in a single bin. By setting the first FFT bandwidth to 800 Hz and using a signal that lifts the main spectrum by less than 1 dB, one can find the expected behavior when analyzing S/N in the baseband with a narrow filter. The noise level at the timf2 input is the "Floor" value. See Fig 2. The "First FFT amplitude" should be set for this value to be about 1.5, 14 dB above the quantization noise, when nothing is connected to the RX144 input. For a system noise figure of 0.4 dB, using the assumptions of Table 2, the "Floor" value will grow to about 5.7 when the preamplifiers are connected.

In timf2, the reverse FFT in 8192 points, the signal would grow by up to 8192 times or 13 bits if no right shifts were used in the butterfly loops of the reverse FFT routine. A number of the butterfly loops use a right shift to prevent the signal from growing and these right shifts introduce errors, another form of quantization noise. It is important not to set the number of butterfly loops with a right shift larger than necessary to avoid this noise but on the other hand it is important to have as many right shifts as possible to allow large interference pulses in the timf2 output. A continuous carrier, a single large frequency bin, will not cause saturation in a reverse transform. Its large amplitude will not grow, it just spreads out over the entire time spanned by the backward transform. Pulses behave differently. A noise pulse in the frequency domain is spread out over all frequency bins. The back transformation will collect all the energy into a single point in time, causing very large amplitude and possibly an overflow since only 16 bits are used.

Table 3 shows the effect of different values of "First backward FFT *att. N*" with the other parameters as described above. The table shows signal and noise levels when a weak signal is injected into one of the RX144 inputs. Rounding errors cause a small loss of signal and an increased noise floor. The signal level is equal to the noise level in 4 kHz bandwidth, but the levels are measured in 1 Hz bandwidth to provide the 0.1 dB accuracy of the table while the noise is measured in a bandwidth of 1 kHz. The test signal is 22 dB above the noise in the bin bandwidth of the first fft. A strong signal will be less attenuated, but a really weak signal will not be more attenuated. The right shifts are placed as late as possible among the fft butterfly loops and the test signal is already below the noise floor when it becomes attenuated as shown in the table. An inspection of Table 3 indicates that the correct value for "First backward FFT *att. N*" is five. The associated loss of noise figure at the RX144 input is about 0.2 dB.

The 1-dB compression point of the RX144 is about +15 dBm. Pulses that have a peak power of +15 dBm after passing a filter with a bandwidth of 2 MHz reach the input of the RX2500 with a peak power of +3 dBm. The reduced power level is not due to amplitude clipping; it is because of the reduced bandwidth. The output bandwidth of the RX10700 is about 0.5 MHz, so the pulses are stretched by a factor of four with four times less power in each pulse causing a peak power reduction of 16 times (12 dB). The pulses that nearly saturate the RX144 input do not saturate the Delta 44 A/D converter although the margin is only 2 dB. Very large pulses do saturate timf2 to an extent that is determined by the "First backward FFT *att. N*" parameter. Table 3 shows the maximum pulse level at the RX144 input that will not saturate timf2 for different values of the parameter. The data is from measurements with a preamplifier having a bandwidth of 2 MHz.

The last entry of Table 3 is the level in dBm at the RX144 input that will cause saturation at the output of the first reverse transform when a signal

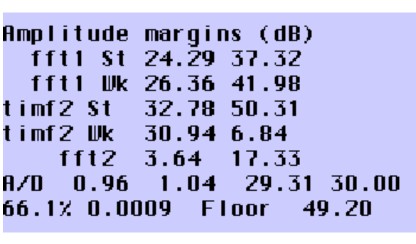

Fig 2—The lower left corner of the Linrad screen when "A" has been pressed. A strong carrier, 3 dB from A/D saturation is fed into channel 1 of the RX144, while –16 dBm interference pulses with a repetition frequency of 100 Hz are fed into channel 2 (see text).

Table 3

The number of butterfly loops with a right shift affects S/N and the saturation level of timf2. The gain levels of earlier stages affect this table, which is for an fft1 size of 8192 with a $\sin^2$ window and with "First FFT amplitude" 1100 to place the noise floor at 1.5 bits RMS with dummy loads at the RX144 input.

Att. N	*Signal (dB)*	*Noise (dB/1 kHz)*	*Max Pulse (dBm)*	*Max Abrupt (dBm)*
2	24.7	18.4	–28	–54
3	24.7	18.4	–22	–48
4	24.7	18.4	–16	–42
5	24.6	18.5	–10	–36
6	24.4	19.0	–4	–30
7	23.8	21.0	+2	–24
8	23.1	25.5	+8	–18

is switched on or off abruptly within the visible passband. An abrupt switching will cause a keying click that spreads its energy over the entire passband. The mechanism is the same as for the interference pulses. This maximum abrupt signal level becomes smaller if the operator selects to use averaged spectra to locate very strong signals. It is not really a big problem because the interference spike created will happen only once for each transmission period. The strong signal must be absent for a few seconds for the gain to go back to normal at the frequency in question.

The strongest interference pulses that will be correctly treated by the smart blanker is –10 dBm. Pulses above this level will be removed by the dumb blanker.

The quantization noise gives rise to spurs, but these spurs are harmless because they disappear when the preamplifier is added, a phenomenon usually referred to as dithering. The amplitudes of the quantization noise spurs are generally independent of the signal level. When a single weak signal is fed into the RX144 input with the above parameters, and only two or three frequency bins are routed to "fft1 St", the group of strong signals, the output, "timf2 St", will be zero most of the time with occasional occurrences of one bit in either direction. This is, of course, no good representation of a sine wave, the signal is surrounded by strong spurs. When preamplifier noise and/or other signals are added, statistics will take care of these spurs.

With the parameters described above, my 600 MHz Pentium III uses 66% of the time available to *Linrad* for computing while spending 34% in the idle loop as can be seen in Fig 2. The idle loop goes to sleep regularly so the *Linux* kernel or other programs may be active in parallel. One cannot be sure all the 34% would be available to *Linrad* if the sleep statement were replaced by useful processing. It may depend on *Linux* activities that I do not know anything about.

The second number on the last line of Fig 2 is 0.0009. This is the longest time in seconds encountered for the idle loop. If the kernel makes lengthy activities due to some other program this number will grow if it happens while *Linrad* is in the idle loop. This number is an indicator for *Linux* doing other tasks than *Linrad*'s signal processing. It will grow while data is saved to disk for example.

The value 49.2 for "Floor" in Fig 2 is the flat noise floor of the pulse train in channel 2. It is about 30 dB above the level of 1.5 with nothing connected to the RX144 inputs. At a repetition frequency of 100 Hz, pulse noise up to 36 dB above the RX144 noise floor will be properly treated by the smart blanker, which means that pulse noise up to about 20 dB above the preamplifier noise floor will be properly handled. This may seem inadequate, but a comparison with the peak power S-meter readings of a conventional radio is irrelevant. For real power-line interference, typically a few thousand pulses per second, the smart blanker will completely eliminate pulses that lift the noise floor by more than 30 dB above the preamplifier noise floor.

The timf2 margins reflect the "First backward FFT *att. N*" setting. Pulses about 30 dB from saturating the A/D converter leave a margin of about 7 dB until saturation occurs in timf2 Wk. It is ok for timf2 Wk to saturate occasionally, but nothing else should saturate. If fft1 or fft2 saturate, strong spurious signals would be generated.

The "Second forward FFT *att. N*" parameter is set to 9 for the result shown in Fig 2. This parameter adjusts the gain of fft2 by selecting how many of the butterfly loops should have a right shift. If this parameter is set too high, quantization noise will add to the noise floor as one can see in Table 4. "Sellim maxlevel," the parameter that controls the maximum permitted amplitude in a single fft1 frequency bin must be set to 4000 or less in order to avoid fft2 saturation for a strong and very stable carrier. Such signals are unlikely in real usage, and if you note fft2 is never near saturation you may make this parameter bigger, which will make the waterfall diagram give a better representation of strong signals.

Summing up, for the WSE RX converters, the following FFT parameters should be close to optimum for 144 MHz EME:

- First FFT bandwidth (Hz) = 30.
- First FFT window (power of sin) = 2.
- First forward FFT version = 5.
- First FFT storage time (s) = 4.
- First FFT amplitude =1100.
- Enable second FFT =1.
- First backward FFT version =1.
- Sellim maxlevel =4000.
- First backward FFT *att. N* =5.
- Second FFT bandwidth factor in powers of 2 =3.
- Second FFT window (power of sin) =1.
- Second forward FFT version =2.
- Second forward FFT *att. N* =9.
- Second FFT storage time (s) =20.

Table 4

The number of right shifted butterfly loops in fft2 affects the noise floor. Parameters are as in Table 3 with "First FFT *att. N*" = 5.

Att. N	*Signal*	*Noise*
6	35.8	–12.4
7	35.8	–12.4
8	35.8	–12.4
9	35.8	–12.4
10	35.8	–12.2
11	35.8	–11.6
12	35.8	–9.9
13	35.8	–6.0
14	35.7	–0.9

Linrad Setup: AFC, Spurs and Baseband

When AFC is enabled, the user must supply parameters that determine how much memory will be allocated. One of these parameters is "Second FFT storage time (s)," for which 20 seconds is a reasonable value. EME signals on 144 MHz are fairly stable, the default values "AFC lock range Hz" = 150 and "AFC max drift Hz/minute" = 100 should be perfectly adequate. Do not enable Morse decoding, those routines are experimental and will not be useful in the near future.

The spur-removal algorithm uses the same spectra as those used by the AFC. The AFC needs high resolution for optimum sensitivity and that is the reason the fft2 bandwidth is set to 2 Hz with the parameters suggested above. The spur removal works like a PLL that sets up a sine wave with the correct amplitude and phase to match the amplitude and phase found in the fft2 transforms over some time selected by the user. The minimum number of transforms is three, the spur-cancellation PLL will fail if the bandwidth of a spur is above 0.2 Hz or so with the above parameters. The spur-removal routine can lock to a peak in the fft 2 spectrum and remove it only if it is coherent from transform to transform. This means that only spurs that are narrow with respect to a 2 Hz bandwidth will be removed. Set "Max no of spurs to cancel" to 100 and make "Spur time constant (0.1sek)" equal to 1.

The maximum bandwidth one would ever want when listening to an EME signal is 2 kHz, which means that the baseband sampling speed should be set to at least 4 kHz. The baseband is filtered out from the fft2 spectra and the total spectrum width must be about 4 kHz for a flat region of 2 kHz. The baseband sampling speed must be a power of two smaller than the input sampling speed so the desired value for "First mixer bandwidth reduction in powers of 2" is four, which leads to a baseband sampling speed of 6857 Hz and a maximum bandwidth of 3.0 kHz. On my 600 MHz Pentium III , the largest baseband transform, fft3, that can be used is 16384 at this relatively high baseband sampling speed. That means that the largest usable baseband filter spans 2.4 seconds so the narrowest carrier filter that can be used for coherent CW will be about 0.5 Hz. This is perfectly adequate for EME but for low bands one may select a much lower baseband sampling rate for coherent CW at very slow speeds.

The "First mixer no. of channels" must be set to one. Some day, when the Morse decoding routines are in place, it will be possible to have the CW transmissions of several stations decoded simultaneously on the screen. The idea is to be able to see what other stations do while operating. This should be very useful in contests for example. The "Baseband storage time (s)" is mainly for Morse decoding. Set it to 20 seconds to not waste memory. When you select 3 kHz bandwidth, the baseband storage will then need 13 MB, but for CW reception with a bandwidth of 20 Hz the memory needed will be 200 kB only. The baseband power spectrum can then be averaged over 20 seconds maximum but that is sufficient for EME CW.

The "Output delay margin (0.1sek)" parameter adds an extra delay between input and output to allow for the computing delay. On my computer, three is enough here. When this parameter is set too low, there will be gaps in the output signal occasionally when the computed data is not available in time for the output. Press "T" on the main screen to see the timing information. The line "D/A" shows the current value and the minimum value encountered. If the minimum becomes zero, the delay margin is set too small or the computer is doing other tasks that slow down processing temporarily. There is no reason to set "Output sampling speed (Hz)" above 6000. High speeds here cost a lot of CPU time because I have not optimized the code for that. The baseband data is present at a sampling rate corresponding to the bandwidth of the baseband filter. For a 20-Hz baseband filter bandwidth, the baseband sampling rate for timf4, is only 47 Hz. To convert this to the desired output frequency, Lagrange's interpolation formula is used to interpolate each output point from four baseband data points, a third-order polynomial fit. The reason is that the output may be on a different sound card with a noninteger ratio between input and output sampling speeds. The procedure is efficient to convert between similar sampling speeds that are related by fractional numbers when the signal is not over-sampled. Four terms are then needed to avoid introducing distortion. When the output sampling speed is set to 96 kHz, this procedure becomes ridiculously inefficient. I see no reason to provide a routine for converting a narrowbandwidth signal to a high sampling rate.

The output mode is a number that characterizes the baseband processing. This number changes when you click on the different boxes in the baseband graph. The current value is shown in the lower right corner of the baseband graph. Set "Default output mode" to the number you want as the default mode. The last parameter "Audio expander exponent" is the exponent by which the amplitude is expanded when the operator clicks the "Exp" box. Expanding the audio volume may be helpful when a very narrow bandwidth is selected. The ears have a logarithmic response for amplitudes. When a matched filter that will only let through the signal and the principal sidebands is used, the ears will have to rely on amplitude information only because the human hearing system does not have the selectivity to distinguish different frequencies within a 15 or 20 Hz wide passband. It then helps to expand the dynamic range of the audio signal. The default value is three.

Summing up, the optimum AFC, spur and baseband parameters for 144 MHz EME should be something like this:

- Enable AFC/SPUR/DECODE = 1.
- AFC lock range Hz = 150.
- AFC max drift Hz/minute = 100.
- Enable Morse decoding = 0.
- Max no of spurs to cancel = 100.
- Spur time constant (0.1sek) = 1.
- First mixer bandwidth reduction in powers of 2 = 4.
- First mixer no of channels = 1.
- Output delay margin (0.1sek) = 3.
- Output sampling speed (Hz) = 6000.
- Default output mode = 1.
- Audio expander exponent = 3.

Receiving a Weak EME CW Signal

With the parameters listed above, the waterfall graph is very sensitive. The FFT size is 65536, but the screen is only 1024 points on my computer. Consequently each pixel on the screen represents 64 frequency bins of the fft2 spectra.

Rather than showing the average power over 64 frequency bins, which would produce the same result as an average over 64 transforms of size 1024, each pixel on the screen shows the strongest frequency bin out of the 64 behind each pixel. This becomes particularly favorable when the fft2 spectra are averaged before the strongest frequency bin is picked.

Setting "Waterfall avg" to six will give a new line on the waterfall every three seconds with a sensitivity that will allow the operator to see all signals present on a 90 kHz segment of the 144 MHz band well below what will be possible to copy. A one-minute transmission is well visible if the S/N ratio is –6 dB in 20 Hz bandwidth. To copy Morse code, one needs something like 14 dB more. Taking the effects of fading into account, copying is done during a few signal peaks when a few letters are above the threshold and while the average signal is at S/N close to zero. When the waterfall graph is expanded to show 1/64 of the spectrum only, the sensitivity is about 3 dB better. Picking the best peak rather than computing the average is an advantage of about 6 dB with the above parameters.

The waterfall graph of *Linrad* shows the total power spectrum summed over both polarizations when a crossed-Yagi array is used. Compared to a perfectly aligned antenna, this means a loss of 3 dB in detection sensitivity. It is not a simple sum of two power spectra because that would lead to an even greater loss in case the polarization is not aligned to one or the other antenna. For each frequency bin, the power of each channel is averaged separately and the complex correlation between the two amplitudes is also averaged. A signal that is present in both channels simultaneously will produce a non-zero average correlation, which is taken into account when computing the energy content of a frequency bin. This is necessary to have a good sensitivity for signals that have a polarization that puts about 50% of the power in each channel.

With the parameters listed above, the minimum processing delay is four seconds when the AFC delay is set to zero. For extremely weak signals delays up to about 10 seconds may be useful. The operation does not differ from the operation described earlier.[3] When the mouse is clicked on a signal, the two channels are analyzed and the polarization is extracted. Depending on the operator's preferences, the two channels can be combined to two new orthogonal polarizations, one has all the signal energy or they can be both routed to stereo headphones.

The EME window, Fig 3, uses the polarization of the received signal to calculate the optimal transmit polarization. This way the adverse effects of Faraday rotation can be eliminated both for receive and transmit. The EME window shows the moon position for this location and for a DX location. A call sign, or fragments thereof, can be entered in the largest box. Fragments must be separated by question marks or stars to indicate one or many unknown characters. Typing in V?2F* will hit VK2FLR as the only answer. VK2* will suggest three call signs while *2FL* will suggest JO2FLD besides VK2FLR and V*LR will suggest VE6LR and VK2FLR. The EME database files dir.skd, eme.dta and allcalls.dta can all be downloaded from the Internet. The EME installation procedure will search them all and collect inconsistencies in a file, while creating a text file containing call signs and locations only. The text file can be loaded automatically when *Linrad* is started.

The Future

In my experience, more bandwidth is more important than anything else. An analog noise-blanker that operates at a bandwidth of 5 MHz is capable of removing very strong static rain noise. S9 noise that sounds exactly like normal white noise can be completely removed. I think the *Linrad* blanker will do it at much lower bandwidth than 5 MHz but 90 kHz is most probably not enough. I have not had any opportunity to make a test, I am still without an antenna since a big storm two years ago.

In the future, when the “standard PC” has a lot of unused CPU power when processing a 0.5-MHz bandwidth, one can make significant improvements to *Linrad*. An improved process could look like this:

1. Forward fft from raw data
2. Back transform for weak signals only
3. Smart blanker subtracts pulses from raw data and remembers what was subtracted
4. New forward fft from improved raw data
5. Strong signals of known types are analyzed. It is possible to model the nonlinearities of, for example, an SSB transmitter and calculate the signal components over the entire spectrum. Known signals are subtracted from the improved raw data to produce new raw data with much lower signal and interference levels. What is subtracted is remembered for further use.
6. New forward fft from better improved raw data
7. Back transform for weak signals only
8. Add the pulses that were subtracted in step 3 and run the smart blanker again. This time the pulse shapes will be very accurate. They are removed from the original raw data.
9. Refine the strong signals and remove them.

The basic idea is to split the total input signal into a few groups of accurately known signals for which *Linrad* can calculate the true waveform based on knowledge of the signal source. The operator can select one of these signals or use a receiver that operates on whatever remains when the strong signals are subtracted.

For the HF bands, a very large bandwidth is probably not so useful. A large number of channels on the other hand would be extremely useful, since *Linrad* could then form an adaptive antenna that optimizes the pattern for optimum S/N for each interference source. In that way it will be possible to overcome very large interference levels from all the modern electronics and so on. With many channels, it will be sufficient with a mediocre dynamic range for each channel so simple systems sampling directly from the antennas would be adequate. We just have to wait for the hardware cost to become low enough.

Comparing the WSE Converters to Conventional Receivers

Blocking Dynamic Range

Blocking dynamic range, BDR is

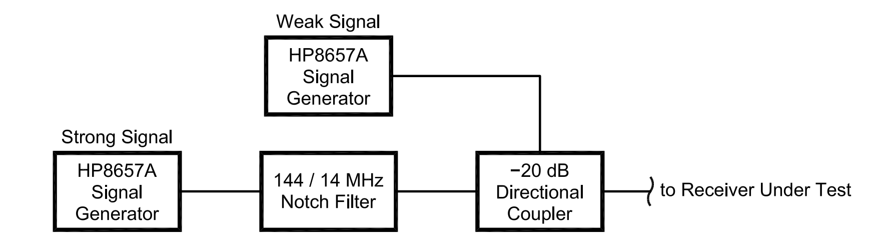

Fig 4—The strong signal passes a notch filter that removes the phase noise from the HP 8657 at a fixed frequency. A weak signal at the notch frequency is injected through a directional coupler towards the receiver under test while the strong signal is picked up by the directional coupler to allow a precise determination of the level entering the test object.

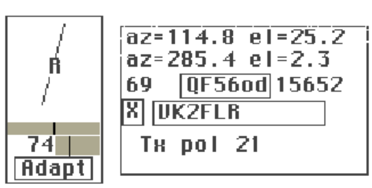

Fig 3—The polarization graph, left and the EME graph right. At this moment, the signal from VK2FLR was received in a nearly vertical polarization. The optimum transmit polarization is 21°. When using H for transmit, the loss due to misalignment is 0.6 dB, but when using V for transmit, the loss is 9 dB. Knowing what to choose improves the QSO chance by a factor of two in this case. The direction to VK2FLR is 69° and the distance is 15,652 km. Direction and distance are intended for terrestrial work. It is possible to enter a locator in the locator field.

Table 5

Blocking and BDR for WSE RX144 system and for a IC-706MKIIG on 144 MHz. N indicates abrupt increase of noise floor due to op-amp saturation in RX2500. The preamplifier is off for the IC-706.

	***WSE*+Linrad** Level for		***IC-706 144 MHz*** Level for	
Frequency offset (kHz)	*3 dB S/N loss (dBHz)*	*1 dB sat (dBHz)*	*3 dB S/N loss (dBHz)*	*1 dB sat (dBHz)*
5	145	A/D sat	102	119
10	145	A/D sat	107	133
20	150	151N	116	139
30	162	163N	120	142
40	164	165N	123	145
50	166	168N	125	146
100	167	172	131	146
250	171	173	133	146

defined in words as: "The ratio (difference in dB) between the weakest on-channel signal a receiver can hear and the strongest off-channel signal a receiver can tolerate without degradation of the received signal." Notice that this is quite different from BDR as measured by ARRL Lab. They measure the level at which blocking occurs.

To measure the dynamic range properly, one needs a strong signal of extreme purity and a weak signal that is not critical. To demonstrate the performance of the WSE converters used together with a modified Delta 44 sound card, I have made the BDR measurements shown in Tables 5 and 6. The measurements were made with the setup shown in Fig 4. Table 5 shows a comparison of the RX144 in a late prototype stage together with production units of the RX70, RX10700 and RX2500.

The RXHFA converter was in a very early prototype stage when this was written. The system noise figure of the entire 14 MHz receiver with the RXHFA prototype operated together with the RX70, RX10700, RX2500 units and a Delta44 in minimum gain mode is 17 dB. A comparison between Tables 5 and 6 shows that the local oscillator of the RXHFA prototype needs some further improvements. This oscillator must operate at several well-separated frequencies to cover amateur bands from 1.8 to 14 MHz, the LO buffer amplifier is the dominating noise source.

The weak signal was set to a level of about 10 dB above the noise floor and the level at which the strong signal degrades S/N by 3 dB was located at several frequencies. The AGC of the transceivers was not switched off, AGC makes no difference because both signal and noise were monitored with *Linrad* running as an audio spectrum analyzer. Table 6 shows that the Japanese transceivers are limited by reciprocal mixing and that *Linrad* and the WSE converters can tolerate about 20 dB higher interference levels. The RXHFA unit may need an attenuator to shift the A/D saturation level upwards in case peak powers above –12 dBm are encountered within the 90 kHz passband. The FT-1000D can receive such signals without an attenuator, but S/N would be degraded seriously by reciprocal mixing so the RXHFA unit will perform better even with an attenuator in front of it. Note that the IC-706 is better than the FT-1000D in case the interference is within ±25 kHz because the LO phase noise is lower.

In the BDR test, I have chosen to measure the level at which S/N is degraded by 3 dB. There is a good reason for selecting this rather than the 1-dB degradation point, which would be more conventional. The time for the measurement increases drastically, or the accuracy is degraded, if one looks for the point of 1-dB degradation. Noise adds by power, converted to a decibel scale it looks like Table 7.

If one wants to determine the level of the added noise within ±1 dB, one must measure a 3 dB change within ±0.5 dB, but one would need to measure a 1 dB change within ±0.2 dB, something that would require a 6.25 times longer integration time when measuring the noise floor.

For use on crowded HF bands, it might be useful to measure the level of the strong signal required for say 15 and 30 dB S/N degradation. In some receivers, the 1 dB and the 30 dB degradation points are very close, maybe 1 dB apart, while in others they may be separated by up to 35 dB. A saturated A/D converter as well as several other saturation processes cause a highly nonlinear interference growth while reciprocal mixing has a nicely linear behavior. A good operator will know how to insert an attenuator between the antenna and the receiver—or to use the built-in attenuator properly. The attenuator insertion could be automated as suggested by Ulrich Rohde.[4] Personally, I prefer to take such decisions myself depending on the circumstances, but adding a circuit like Ulrich's (in his figure 43) to the WSE converters would be trivial.

The dynamic-range data of Table 6 can be converted from dBHz to dB in 500 Hz bandwidth by subtracting 27 dB. At 20-kHz frequency separation, the result is 100 dB for the IC-706 while it is 97 dB for the FT-1000D. These values represent the true dynamic range in a weak signal usage of the receivers. This is the natural concept to me with a bias from the

Table 6

Blocking and BDR for a WSE RXHFA prototype system, an IC-706MKIIG and a FT-1000D on 14 MHz. G+ indicates that the gain increases rather than decreases when the interference is added. N indicates abrupt increase of noise floor due to op-amp saturation in RX2500.The preamplifier is off for IC-706 and the Front End switch is in position IP0 for the FT-1000D. For FT-1000D blocking is measured indirectly through the cross-modulation from an AM modulated carrier

	***WSE+*Linrad**		***IC-706 14 MHz***		***FT-1000 14 MHz***	
	Level for		***Level for***		***Level for***	
Freq offset (kHz)	*3 dB S/N loss (dBHz)*	*1 dB sat (dBHz)*	*3 dB S/N loss (dBHz)*	*1 dB sat (dBHz)*	*3 dB S/N loss (dBHz)*	*1 dB sat (dBHz)*
5	145	A/D sat	114	122	113	149
10	146	A/D sat	122	129	116	156
20	145	A/D sat	127	143	124	163
30	145	A/D sat	131	149	129	165
40	153	163N	134	149	132	166
50	156	165N	135	150	135	166
100	156	170	140	G+	144	168
250	164	171	148	G+	155	169
500	171	172	149	G+	155	170

Table 7

Adding a second noise source increases the noise level like this. If both noise levels are equal the sum is 3 dB above a single signal and the sensitivity is 0.5 dB for 1 dB change of the added signal. If the added noise is 6 dB below the original noise, the sum is 1 dB above the original noise but the sensitivity is only 0.2 dB for 1 dB change of the added signal.

Added Signal Relative to First Signal (dB)	*Signal Level Change* (dB)
–7	0.79
–6	0.97
–5	1.19
–4	1.46
–3	1.76
–2	2.12
–1	2.54
0	3.01
1	3.54

144 MHz weak-signal community. HF operators may find the distance from the noise floor up to blocking more relevant. Then the ARRL lab procedure might be more relevant. The two methods give numbers that differ by 60 dB! Knowing what the numbers really mean is essential when deciding which radio to buy.

Third-Order Intermodulation

Third order intermodulation, IM3, is typically the phenomenon that limits the dynamic-range performance of a receiver when BDR is not the limiting factor. IM3 can be described as frequency mixing due to the nonlinearity of amplifier, mixer or other stages that arises when the signal levels are very high. When two signals f1 and f2 enter a receiver, IM3 is produced at frequencies that can be described as the difference between one signal and the overtone of the other signal, 2f1–f2 for example. In this case, the IM3 level is proportional the f2 and to the square of the f1 signal levels. In a two-tone test with equal amplitudes for f1 and f2 , the IM3 level is proportional to the common signal level to the power of three. This is the third-order law saying that for a 1 dB increase of the signal levels the IM3 levels will increase by 3 dB. The third-order law is the basis for this definition: *The third order intercept point (IP3) is the point at which the the extrapolated third-order intermodulation level (IM3) is equal to the signal levels in the output of a two-tone test when the extrapolation is made from a point at which and below the third-order intermodulation follows the third-order law.*

There are different procedures suggested for the measurement of IP3. How to make the measurement on a mixer or preamplifier is uncontroversial, but how to handle a "black box" with an antenna input and a loudspeaker output is less clear. Some receivers have an AGC that cannot be switched off, and there may be other complications. Procedures to measure IP3 may give a result that is inconsistent with procedures that measure two-tone, third-order intermodulation dynamic range, IM3DR, despite the fact that these two measurements should have an exact relation. They are coupled through bandwidth and noise figure by the third-order law in the relation IP3 = 1.5×IM3DR+NOISE FLOOR. A receiver that does not follow the third-order law [5, 6, 7, 8, 9] cannot be characterized by an IP3 number. The references show a discrepancy of more than 10 dB in the IP3 relation and indicate design inadequacies or measurement errors. I have tried to reproduce the peculiar response reported in Note 5, but found nothing but normal third-order behavior. The TS-450S I looked at had a much later serial number than the one tested in the ARRL Lab and some design inadequacy may have been corrected by the maker in later production units.

There is a simple way to measure third-order intermodulation that will give accurate results regardless of the receiver architecture. It works equally well with AGC on or off and it is very easy to perform. Just combine two equally strong signals and a third, weak one. The IM3 product and the weak signal are placed something like 10 to 100 Hz apart and a spectrum analyzer (*Linrad* for example) is connected to the loudspeaker output. The weak signal is set to give the same amplitude as the IM3 product on the screen. This measurement is fast, easy, reproducible and accurate. The true power levels of the strong signals and of the weak signal that gives an equally strong signal as the IM3 product are measured directly. AGC or AF saturation does not matter. The point of equal amplitudes is independent of the nonlinearities in the stages following the filters that exclude the strong signals. At large frequency separations, a notch filter is useful, just replace the strong signal in Fig 4 by a pair of strong signals that have a frequency relationship that places a third-order intermodulation product at the frequency of the notch. Notice that the quartz crystals in the notch filter produce IM3 at close frequency separations and that a second measurement with an attenuator at the receiver input will show if this is a limitation of the measurement. For measurements at close frequency separations, where a notch filter is useless, the third generator and the audio spectrum analyzer are essential. This is so because the noise and spurs in the two strong signals as well as in the local oscillator of the test object easily lead to incorrect measurements at low IM3 levels.

Real receivers may have peculiarities that make them deviate from the third-order law that is accurately valid for a simple chain of amplifiers and mixers. The reason may be nonlinearities in circuits that are not in the signal path. The noise blanker may have AGC controlled amplifiers that produce modest levels of intermodulation more or less independently of the input signal level. At low signal levels where the intermodulation produced in the signal path is very low, inadequate screening or buffering may allow IM3 from such side paths to interfere with the desired signal. Look at **antennspecialisten.se/~sm5bsz/dynrange/intermod.htm** for a discussion of IM3 measurements, theory, spectra and time-domain waveforms. The site also contains details of the measurements behind the IP3 values presented in Table 8.

For both FT-1000 and IC-706MKIIG, IP3 and IM3DR are degraded by a very small amount if the frequency separation is reduced to 20 kHz from 100 kHz. For the RXHFA unit it is quite different. The bit errors in the A/D conversion process give rise to IM3 that is varying in a seemingly random fashion with the level of the two test tones. The IM3 from the A/D conversion process is at about –140 dBm, below MDS in 500 Hz bandwidth, but it is there. This kind of intermodulation disappears completely if other signals are present in the pass-band as will practically always be the case in the real usage of a receiver. Fig 5 shows the IM3 response of the RXHFA unit for two signals within the A/D converter passband.

As can be seen from Fig 5, the close range IM3 is at the 500 Hz MDS level for a two-tone input of –29 dBm, which means that IM3DR is 101 dB. The A/D converter in the Delta 44 saturates when the levels in the two-tone test are

Table 8

Two-tone third-order intermodulation data at 14 MHz and 100-kHz frequency separation.

Receiver Type	*IP3 absolute (dBm)*	*NF (dB)*	*IP3 to noise floor (dBHz)*	*MDS at 500 Hz bw (dBm)*	*IM3DR in 500 Hz bw (dB)*
RXHFA	25	17	182	–130	103
FT–1000	22	21	175	–126	99
IC–706MKIIG	–4	12	156	–135	86

set to about –18 dBm. The digital output is limited by the number of bits and can simply not represent an analog signal outside the digital range. In the range –30 to –18 dBm, the RXHFA/*Linrad* system follows the third-order law, but it is not fair to characterize the system with the IP3 of +20 dBm one can get from an extrapolation. The RXHFA/*Linrad* system will behave as if it had an IP3 in the order of –8 dBm for multiple input signals that reach a peak power above –18 dBm within the 95 kHz bandwidth seen by the A/D converter. Looking only at the intermodulation, one would conclude that the FT-1000 would be much better in such cases, but the FT-1000 front end will see much higher peak powers because it must handle much more bandwidth. More importantly, the FT-1000 will be limited by reciprocal mixing, the noise floor is degraded by 3 dB at an average signal level of about –30 dBm already. Knowing this fundamental difference between analog and digital receivers is very important. If the bandwidth seen by the A/D converter is even wider or if the dynamic range is lower, IP3 values may still be impressive, but the real intermodulation resistance may be poor compared to “good old analog receivers” with similar IP3 and IM3DR numbers.

The intermodulation characteristics of the WSE converter chain are the same for all frequency bands. The RX144 and the RXHFA units have the same IM3DR. The RX144 is definitely intended to be used with amplifiers in front of it and it will have a noise figure of about 11 dB, which means that IP3 will be about +19 dBm. I have not yet decided whether it is a good idea to incorporate an RF amplifier in the RXHFA to shift the levels downwards. The data given above is without any RF amplifier in the RXHFA prototype and it is compared to the FT-1000D and the IC-706MKIIG with the RF amplifier disabled.

How Much Dynamic Range do We Need?

On the HF bands, the answer is 100 dB for BDR in 3 kHz bandwidth according to Chadwick.[10] This is equivalent to 135 dBHz, which is met easily by the WSE converter chain at all frequency separations, but which is also met by IC-706MKIIG and FT-1000D at frequency separations above 50 kHz. As I read the referent of Note 10, this is good enough on the HF bands. On 7 MHz one may need an IP3 of +36 dBm at a noise figure of 33 dB, which is just about what the FT-1000D can perform with 12 dB attenuation, but which is met with some margin by the RXHFA unit when a 15 dB attenuator is added. At other times, a noise figure of 22 dB is needed. The operator must be able to move the dynamic range levels up and down with an attenuator, but with that constraint, modern receivers are good enough for the HF bands.

On 144 MHz, my favorite band, it is quite different. Fellow amateurs typically cause the most difficult problems. A 2 m station may put 100 W

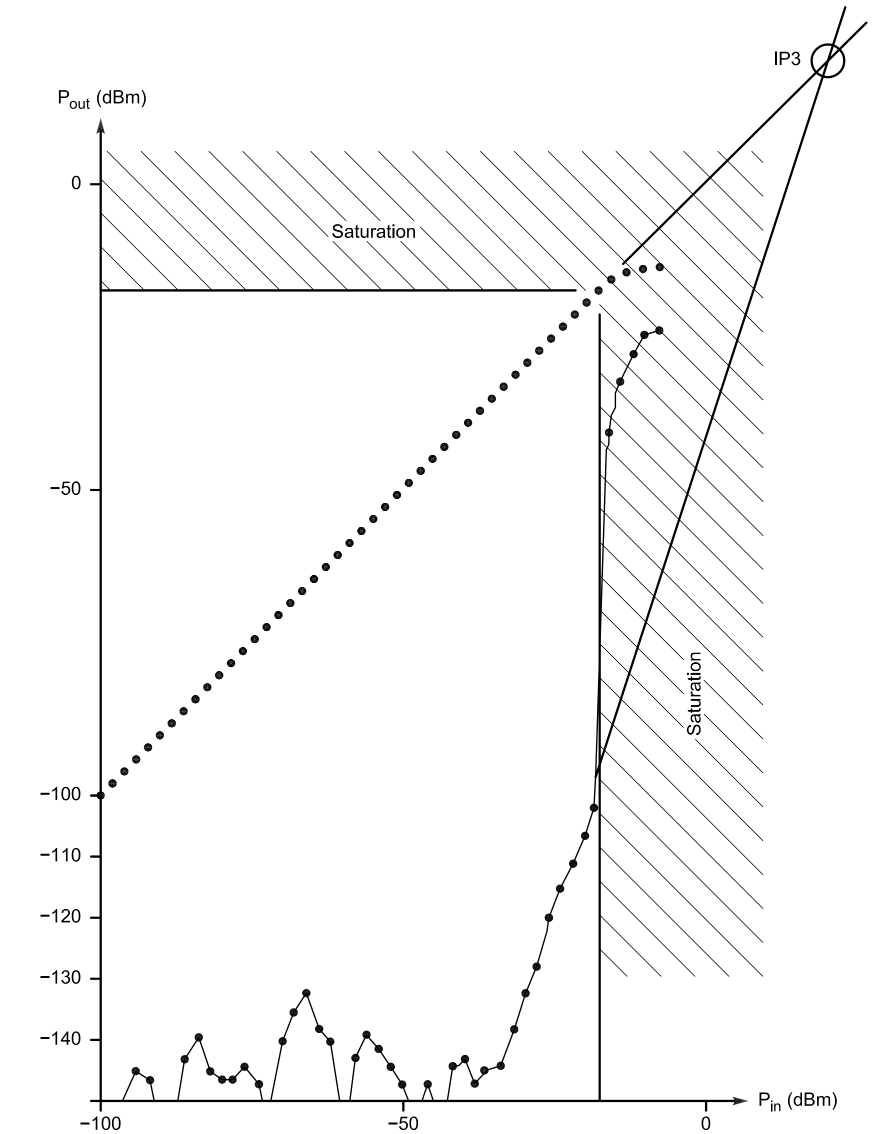

Fig 5—IM3 response for the RXHFA unit in a two-tone test. Notice that the IM3 below –35 dBm is real but that it disappears due to dithering if noise or another signal is added.

Table 9

Received power levels with antennas pointing into each other on 144 MHz, at two different distances assuming free space propagation.

Distance (km)	*Tx ERP (kW)*	*Rx Power Ant = 13 dBd (dBm)*	*Rx Power Ant = 18 dBd (dBm)*
1	3	+2.4	+7.4
1	100	+17.4	+22.4
10	3	–17.6	–12.6
10	100	–2.6	+2.4

into a 13 dBd antenna to produce an effective radiated power of 3 kW. Much higher ERPs are not uncommon, 1 kW into four modest Yagis will easily give an ERP of 100 kW. Add to these high power levels the much greater receiver sensitivities, and the directional gain at the receiver side and you will find received powers at some different distances as illustrated in Table 9.

Contrary to the HF bands, interference on 144 MHz is likely to be caused by one or very few signals. The reason is the high directivity of the antennas. Having one of the local high power stations pointing his antenna into my direction while I point my antenna in his direction is not likely to happen simultaneously for many local high power stations. This means that BDR is generally more important than IM3DR on 144 MHz. RX144 provides 145 dBHz for close spaced signals and with a noise floor of –174 dBm/Hz it means that the maximum permitted signal level is –29 dBm. Table 9 indicates the need for much higher levels. On 144 MHz, we often run into mutual interference because of inadequate dynamic range. At frequency separations above 50 kHz, the RX144 provides 166 dBHz so the maximum permitted signal level is –8 dBm. Table 9 indicates that much more could be useful sometimes, and –8 dBm does not allow unperturbed reception, it is the level where S/N is degraded by 3 dB. Here the influence of dynamic range loss and noise figure due to the preamplifier as illustrated by Table 2 is neglected, but these effects work in opposite direction and cancel if both are made about 3 dB.

Conventional transceivers often produce strong noise sidebands. I have measured several transceivers using RX144 and *Linrad* as a spectrum analyzer. The data is available at **antennspecialisten.se/~sm5bsz/dynrange/gavelstad/gav.htm**. The noise floor is typically at –110 to –120 dBc/Hz at a frequency separation of 20 kHz and –125 dBc/Hz at 50 kHz. To me the WSE converters and *Linrad* is not only a radio receiver, the system is also an instrument for finding cures to the problems caused by design inadequacies in commercial transceivers. A good LO, such as the one in TM255E is at –137 dBc/Hz at 20 kHz, something that is proven by the excellent BDR value, but the transmitter noise is at –122 dBc/Hz because other noise sources than the LO dominate the transmitted signal. The dynamic range needed to make the WSE converters useful as laboratory instruments is 10 dB better than the best transmitters one would want to investigate. At today's state of the art, the performance is just about good enough in the close range, but as soon as the interference is outside the frequency range routed to the A/D converter the performance is adequate with a good margin.

Notice that noise levels I present are always RMS values. They truly reflect the power ratio between the noise power in a defined bandwidth and the power of a carrier. It is a bad habit among engineers to interpret as dBc/Hz decibel numbers that come from the display of a spectrum analyzer that averages the output from a logarithmic detector. Such decibel values are about 6 dB lower than the true dBc/Hz values in which the ARRL Lab composite noise test is defined.[11]

When *Linrad* and the WSE converters are used to measure sideband noise, the numbers obtained are about 6 dB worse than those published in *QST* because *Linrad* computes the true RMS power levels. The –145 dBc/Hz noise floor within the frequency range seen by the A/D converter therefore corresponds to –151 dB in the ARRL Lab scale.

Conclusions

It is demonstrated above that the WSE converters and *Linrad* give a third-order dynamic range that is comparable to good analog receivers while the BDR is much better. The data is based on measurements on prototypes, but the final outcome will not be very different.

Linrad is not designed for the WSE converters, it is intended to be used in the future with very much simpler digital hardware that makes the A/D conversion at VHF frequencies and samples the antenna signal directly. I have designed the WSE converters because it was reasonably simple with the tools at my disposal, and I did not want to wait for someone else to produce the digital hardware and drive routines for *Linux*. Another reason is that I believed it was a way to get a performance that is somewhat better than I can expect to ever get from a digital system. The WSE converters will be the radio I use in the future, but they will also constitute the tools needed to verify the operation of the digital hardware when it becomes available. The digital revolution will continue. As amateurs, we face a new and exciting situation in which we can take a leading role in the development of new technologies. By feeding more bandwidth from more antennas into a computer it will be possible to remove interference to an extent we would not even dream of today. Imagine 16 ferrite rods that are placed around your location sending digital data to your computer, each one with a battery, a small digital processing block and a microwave link. The battery could be powered by solar cells. The computer can form an adaptive antenna with 12 dB gain for each interference source, then subtract the interference with a very high accuracy, if the interference has any characteristics that the computer can be programmed to identify. Finally, the adaptive lobe can be pointed towards the desired signal, which will become readable even if it is deep below the interference level in a single antenna. Personally, I think the strategies to identify and remove interference form the most exciting field for amateur development in the future. As amateurs, we might want to push the limits in a difficult interference situation on a particular frequency band while a professional would use another frequency or even another technology to avoid the problem. *Linrad* is an early attempt to get into this new field of qualified signal processing, it is *not* just a DSP package for EME enthusiasts.

Notes

[1] L. Åsbrink, SM5BSZ, "*Linrad*: New Possibillities for the Communications Experimenter, Part 2," *QEX*, Jan/Feb 2003, p 41-48.

[2] L. Åsbrink, SM5BSZ, "*Linrad*: New Possibillities for the Communications Experimenter, Part 3," *QEX*, May/June 2003, p 36-43, Fig 1.

[3] L. Åsbrink, SM5BSZ, "*Linrad*: New Possibillities for the Communications Experimenter, Part 4," *QEX*, Jul/Aug 2003, pp 29-37.

[4] U. Rohde, KA2WEU, "Theory of Intermodulation and Reciprocal Mixing: Practice, Definitions and Measurements in Devices and Systems Part 2," *QEX*, Jan/Feb 2003, pp 21-31.

[5] J. Healy, NJ2L, "Product Review: Kenwood TS-450S and TS-690S Transceivers," *QST*, Apr 1992, pp 67-69.

[6] P. Danzer, N1II, "Product Review: Ten-Tec Pegasus HF Transceiver," *QST*, Feb 2000, pp 63-67.

[7] M. Wilson, K1RO, "Product Review: Watkins-Johnson HF-1000 General-Coverage Receiver," *QST*, Dec 1994, pp 76-79.

[8] R. Lindquist, N1RL, "Product Review: SGC SG-2020 Transceiver," *QST*, Oct 1998, pp 71-74.

[9] Rick Lindquist, Product Review: Yaesu MARK-V FT-1000MP HF Transceiver, QST Nov. 2000, p 64 – 69.

[10] P. Chadwick, G3RZP, "HF Receiver Dynamic Range: How Much Do We Need?" *QEX,* May/June 2002, pp 36-41.

[11] M. Gruber, WA1SVF, "Improved Transmitted Composite-Noise Data Presentation," *QST,* Feb 1995, p 62.