

Observations MFSK16 Interference Performance

Jack R. Smith K8ZOA¹

13 January 2001

Summary

A series of measurements were taken to characterize the susceptibility of MFSK16 signals to interference, and to determine the effect of receiver performance upon interference rejection. An optimally configured receiver permits as much as 30 dB more interference rejection compared with a single sideband receiver with a bandwidth set for voice reception. MFSK16 has good performance, with little differences noted between the Stream and HamScope software implementations.

Introduction

As a preliminary survey of the interference resistance of MFSK16, four parameters were measured:

1. The dynamic range in the absence of noise or interference.
2. The required signal to noise ratio for Gaussian noise.
3. The required signal to interference ratio for a single tone interferer spaced between 25 Hz and 1000 Hz from the MFSK16 idle tone.
4. The required signal to interference ratio for an MFSK16 interferer with an idle tone spaced between 0 and 700 Hz from the desired MFSK16 idle tone.

To verify the applicability of the single tone interference measurements, supplemental measurements were made using simulated “off the air” interfering signals and are reported in this paper. Results are presented for five combinations of receiver and audio filter bandwidths.

These measurements were made using both Stream v. 0.88a and HamScope v. 1.2.1 software. Although the purpose of these measurements is not to characterize the performance of Stream versus HamScope, some performance differences were observed and are reported in the discussion.

A significant limitation of the measurements presented is that both the signal and interferer are un-faded. In the absence of an HF fading simulator, therefore, appropriate statistical adjustments are necessary to convert these static measurements to the “real world” fading HF environment. Nonetheless, the data presented herein is useful in assessing the interference susceptibility of MFSK16. Since the interleaving properties of MFSK16 reduce the effect of short fades, the static analysis presented herein is of

¹ Jack R. Smith, K8ZOA, 7236 Clifton Road, Clifton, VA 20124-1802. e-mail: jsmith@mnsinc.com

increased validity, compared with, for example, an unprotected data transmission mode, such as conventional Baudot radioteletype (RTTY) as used by radio amateurs.

A subsequent paper in this series will combine these measurements with the transmitted signal characterization reported earlier and will discuss the compatibility of high power / low power operation of MFSK16 signals, as well as the potential for interference to/from MFSK16 and other modes.

This paper assumes the reader is familiar with MFSK16 terminology and the basics of its modulation methodology. Additional information on these matters can be found at the Stream web site, <http://www.qsl.net/z11bpu/MFSK16>

Test Equipment Configuration

A series of MFSK16 test signals were generated using a Texas Instruments Travelmate 5200, 120 MHz Pentium laptop computer and integrated sound card. These test signals were captured and saved to disk using Spectrogram version 5.1.7 by R. S. Horne.² Spectrogram was run on a Gateway Solo 233 MHz Pentium MMX laptop computer. Spectrogram is a “freeware” sound card audio spectrum analyzer program for Microsoft Windows 9x computers. All data was taken in 16 bit input mode, 44 kb/s sample rate. From theoretical considerations, 16-bit data sampling should yield nearly 96 dB of dynamic range. In practice, somewhat less than this theoretical maximum is realizable.

Each test signal consists of a short preamble, 16 lines of the text message “The quick brown fox jumped over the lazy dog’s back 1234567890 times. K8ZOA sends” and a postamble. Each test signal runs about 5 minutes duration.

Test signals were generated for idle tones of 1000 Hz, 1008 Hz, 1016 Hz, 1050 Hz, 1100 Hz, 1150 Hz, 1200 Hz, 1250 Hz, 1300 Hz, 1350 Hz, 1400 Hz, 1500 Hz, 1600 Hz and 1700 Hz. An idle-only 1000 Hz test was also recorded. After capture to disk, an audio CD-ROM of these captured files was created using a Sony Sprespa CD-ROM burner. Two copies of the test CD-ROM were created, with the thought that through using two CD-ROM players, interference between two MFSK16 signals could be simulated.

The first series of tests involved determining the dynamic range of a typical computer running Stream. The test setup used is illustrated below.

² Spectrogram is available at: <http://www.monumental.com/rshorne/gram.html>.

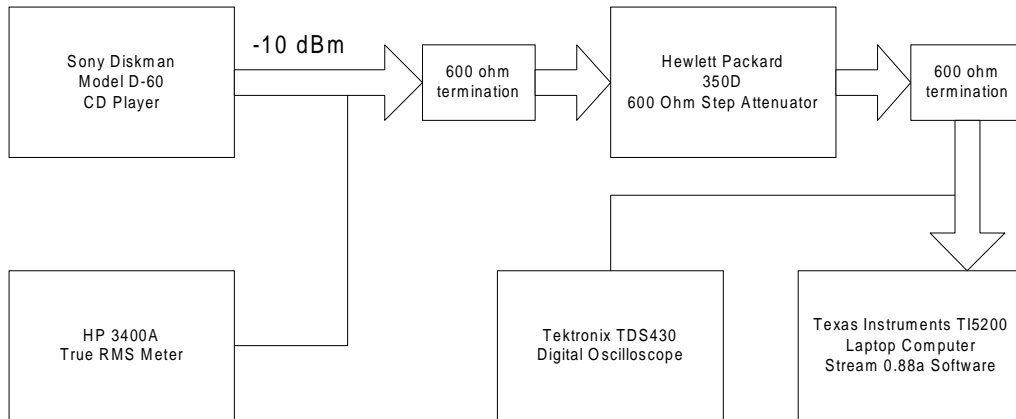


Figure 1 Dynamic Range Test Setup

The Gaussian noise test built on the dynamic range setup, as shown below.

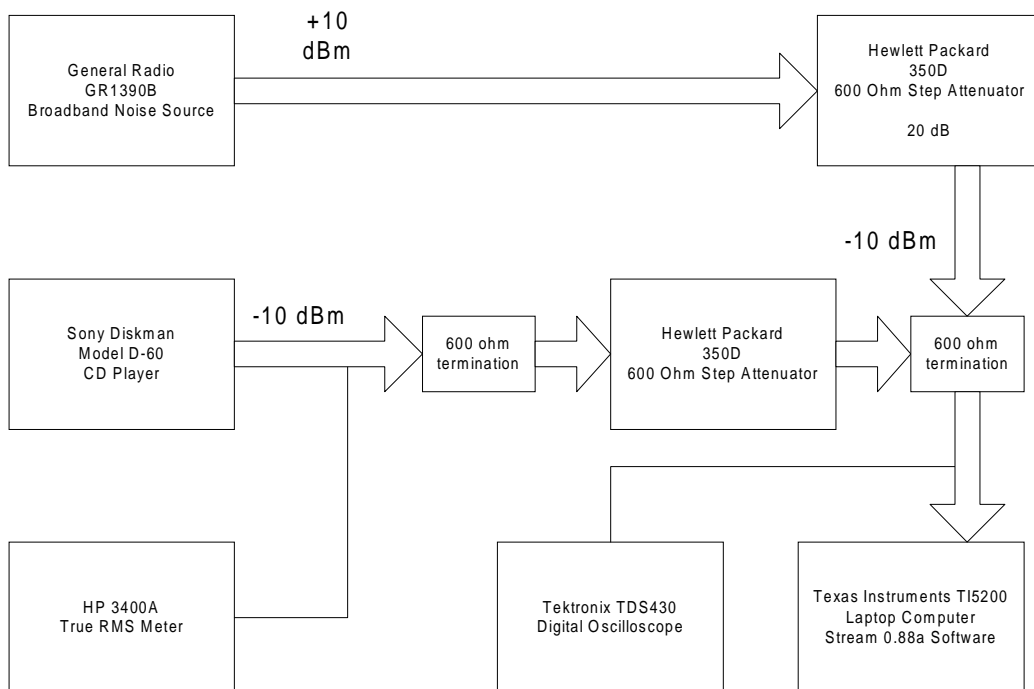


Figure 2 Gaussian Noise Test Setup

The next series of tests reported in this study used a single interfering carrier using the configuration shown below.

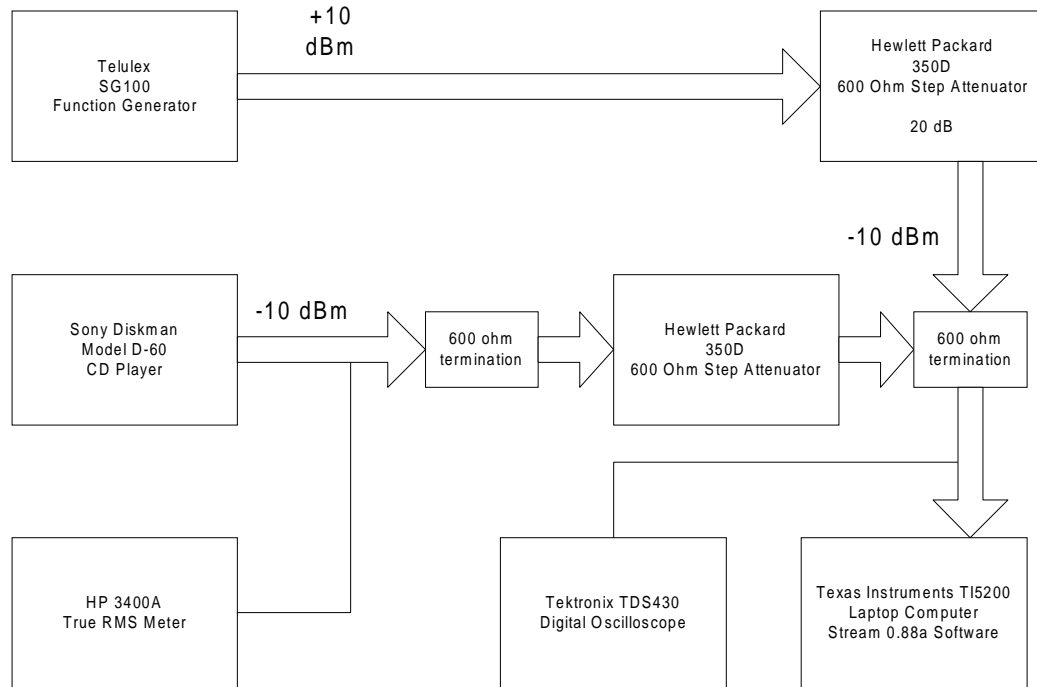


Figure 3 Single Tone Interference Setup

Finally, MFSK16 to MFSK16 interference studies used the test setup of Figure 4.

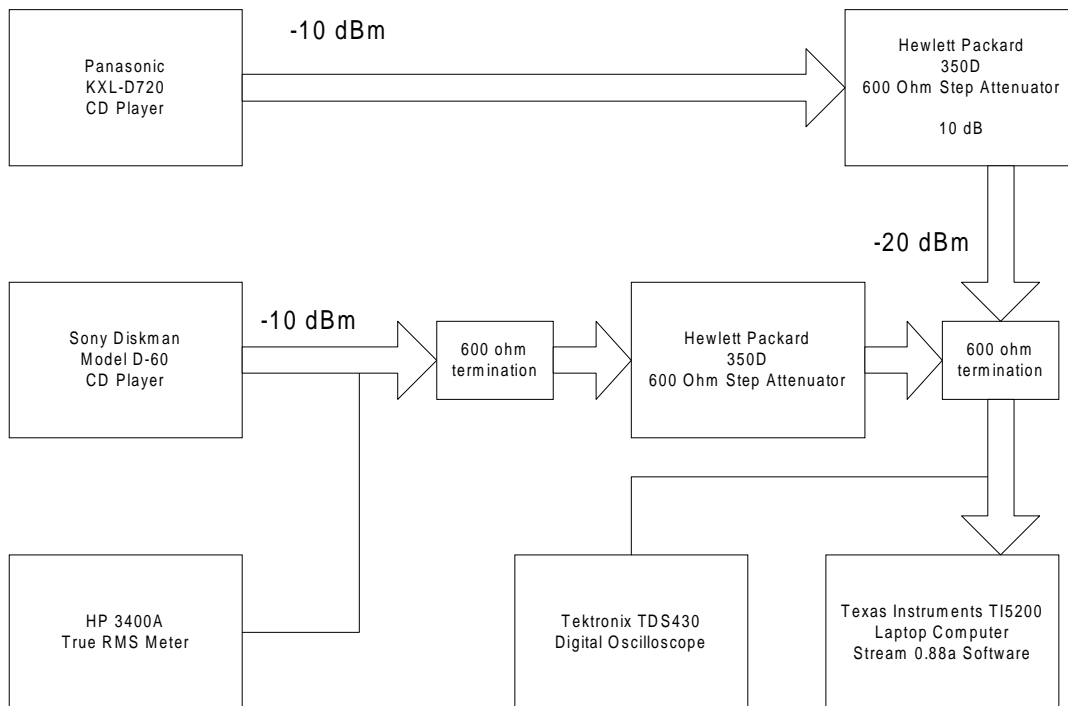


Figure 4 MFSK to MFSK Interference Test Setup

The HP 350D step attenuators permit an adjustment range of 0 to 120 dB, in 1 dB steps. By varying the step attenuator associated with the desired signal (from the Sony D-60 CD Player), various signal to interference ratios can be generated.

The two audio sources were combined without using an audio mixer. In order to keep unwanted mixing products down, a minimum of 20 dB attenuation was used in both legs of the attenuation path, except for the MFSK16-to-MFSK16 tests, where 10 dB minimum attenuation was used.

The test configurations shown above indicates that the TI Laptop computer runs Stream 0.88a. Identical test setups were used for the tests conducted with HamScope v. 1.2.1. In some tests, a Hewlett Packard 200CD audio generator instead of the Telulex SG-100 function generator was used. In these cases, the frequency of the HP200CD was set using a Heath IM-2420 frequency counter.

The Hewlett Packard 3400A true RMS meter was used to set all levels, with either a 600 ohm termination or a 50 ohm termination, as appropriate. Since the 3400A is calibrated in dBm at 600 ohms, an appropriate adjustment was made when setting levels in the 50 ohm environment. Levels were measured out of the generators and the attenuators were spot

checked for proper attenuation. Since all readings are relative, errors in the 3400A are unimportant.

Why aren't there any bit or character error rates shown?

The measurements reported are of the “print or no-print” nature. Normally, the bit error rate, or in the absence of bit errors, the character error rate, is reported in a study of this type.³ However, as coding practices become more robust, and incorporate greater error correction and detection techniques, the nature of the print presented to the user changes. Instead of the gradual degradation seen in, for example, a Baudot RTTY signal, the “cliff” effect is seen. The print is essentially perfect or it is essentially garbage; just like going off a cliff, there are no intermediate points. In the studies made, usually a 1 or 2 dB variation made the difference between perfect print and worthless print. Accordingly, the data presented shows a single “threshold” value. The values presented throughout this report are for “no print.” For “print” an adjustment of 1 or 2 dB is necessary.

The “driving off a cliff” effect is well illustrated with the following sample data, taken from the Gaussian noise performance test with HamScope.

S/N ⁴	Sample Recovered Text	Comment
-4 dB	* * * * Test Starts * * * * * The quick brown fox jumped over the lazy dog's back 1234567890 times. K8ZOA sends The quick brown fox jumped over the lazy dog's back 1234567890 times. K8ZOA sends	Perfect
-5 dB	* * * * Test Starts * * * * * The quick brown fox jumped over the lazy dog's back 1234567890 times. K8ZOA sends The quick brown fox jumped over the lazy dog's back 1234567890 times. K8ZOA sends	Perfect
-6 dB	* * * Test Starts * * * * * The quinrown fox jumped over the lazy dog's back 1234567890 times. K8ZOA sends The quick brown fox jumped over the lazy dog's back 123456n90 times. K8ZOA sends	Very slight garble— useable print
-7 dB	* * TefiJtarts * * * * * Tbe h8nk brown fox jumped over the lazy dogDqou toi 4=a4567890 tbsaA carnWwq dneøeSKlnee enazy dog's back 1234go< tZOA sends	Severe garble-- unusable

³ Stream does not provide a bit error rate display. HamScope displays a Bit Error Rate box, but throughout the tests it remained at 0%.

⁴ S/N is calculated with respect to a notional bandwidth of 316 Hz. See later discussion concerning noise bandwidth. The S/N is negative, which indicates that the desired signal is weaker than the noise power in the 316 Hz bandwidth.

-8 dB	* K k>Ra arawn foi pkEØ rb eddtp l2etnePs. tvIK sen<pMrM aLeabn focnUn bVe irZSXk 123456789Ø times. K8Zo xi sod_,nNikc brei utmped otu e R4'	Unusable
-9 dB	DetgcOt otw *Me ickRkphptaT et ree wed BkeVg k} ØzTte td kfiAl et nocJkijai)XBfs lsi	Unusable

How accurate are these numbers?

The data presented should be regarded as ±2 dB. There is an inherent 1 dB “quantizing” error, since the attenuators have a 1 dB step size. The threshold values are based upon a subjective “good print” or “unusable print” characterization.

Finally, MFSK16 displays hysteresis. For any particular test the desired signal could be decreased from a level providing good print to one of bad print, or vice versa, it could be increased from bad print to good print. The results will be different, by 1 or 2 dB, most likely attributable to synchronization issues. If the measurement starts with a good signal, the software can synch up and the “integrate and dump” filters are properly timed.⁵ This timing will continue for some period after the signal is reduced below the level that would permit acquiring proper timing. Thus the software continues to work in optimum mode. In contrast, if the software has a signal below the acquisition threshold, its integrate and dump timing will be randomly synchronized with the received signal symbol period. As the signal is slowly increased, the software will “snap in” and start properly decoding. In taking the data, the author started with a good signal and decreased the strength. This could be considered optimistic in certain circumstances.

Baseline Observation

Before making error measurements with MFSK16 signals, the 1000 Hz idle test signal was observed using Spectrogram with the TI 5200 computer, with the results shown below. Figure 5 shows the output of the CD-ROM player playing the test recording.

⁵ The author assumes that both Stream and HamScope implement software “integrate and dump” filters. (This fact was confirmed in an exchange of emails with the author of HamScope.) An integrate and dump filter resets the filter output to zero at the precise start of a symbol period, integrates the signal input for the duration of the symbol and then outputs that value. It is then reset for the next symbol. If the start and stop of the integration period is exactly synchronous with the incoming symbol period, the maximum possible energy is extracted from the incoming signal. If it is out-of-phase by one-half the symbol interval, one-half of the possible energy is wasted. In theory, this would amount to a 3 dB synchronization loss. About a 2 dB loss was seen in the measurements taken.

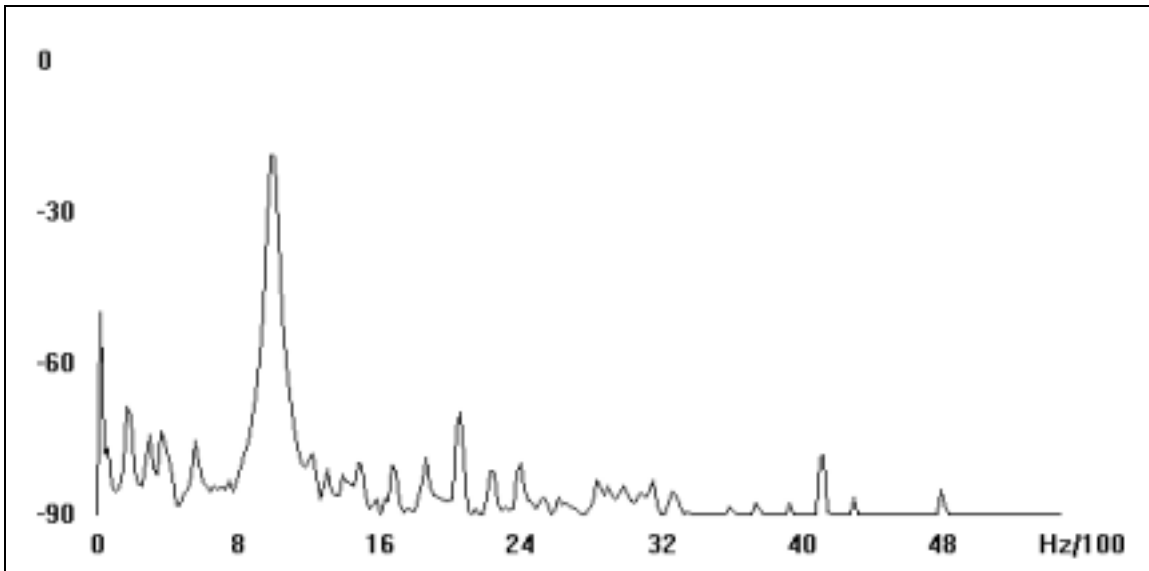


Figure 5 Spectrogram Plot of 1000 Hz MFSK Idle Tone out of Sony CD Player

The desired signal is approximately 60 dB above the extraneous signals.

Dynamic Range

Using the test setup shown at Figure 1, the dynamic range of an unfaded signal was measured.⁶

Adjusting the input signal for -10 dBm input level to the step attenuator, progressively more attenuation was applied until the signal dropped below the threshold of good print. For this and subsequent tests, the audio gain slider in the software driver of the computer sound card was set to maximum gain.⁷

DYNAMIC RANGE	
SOFTWARE	DYNAMIC RANGE
Stream v. 0.88a	85 dB
HamScope v. 1.2.1	86 dB

⁶ The test data was taken using the microphone input to the soundcard. It may be argued that applying a -10 dBm signal to the microphone input overloads the soundcard input stage. To see if this were the case, several measurements were taken using the Line In input, expected to be a good match to a -10 dBm signal. In fact, poorer performance was seen using the Line In input.

⁷ Stream, unlike most other software, does not have an input level indicator. The applied signal level of -10 dBm into a 600-ohm impedance corresponds to 0.245 volts RMS. It is unclear whether this signal level, combined with the maximum gain setting in the soundcard software, is excessive or not. In any case, no print errors were observed when at maximum gain with the full -10 dBm signal applied to the soundcard input. HamScope does have a level indicator, and it did not indicate excessive audio level at the applied -10 dBm signal.

The difference between Stream and HamScope is within the measurement error and should be disregarded. The applied signal in both cases was -10 dBm. Thus, both Stream and HamScope could reliably print audio signals of -94 or -95 dBm.

Gaussian Noise Performance

Using the configuration shown in Figure 2, broadband Gaussian noise was mixed with the test signal and printed by the MFSK16 software. The Gaussian noise level was established at $+10$ dBm into a HP355D attenuator, set for 20 dB attenuation, resulting in -10 dBm of noise at the mixing point.

A brief detour into noise is necessary before continuing further. Gaussian noise is equal power noise, *i.e.*, the noise power in a bandwidth of 1 Hz measured at 1000 Hz is equal to the noise power in the same bandwidth measured at 2000 Hz. The GR 1390B noise source was set to 20 KHz bandwidth, and the output power in that bandwidth was set to $+10$ dBm using a wideband true RMS reading meter, a Hewlett Packard 3400A. Figure 6 shows the noise spectrum is reasonably flat within the specified 20 KHz bandwidth.

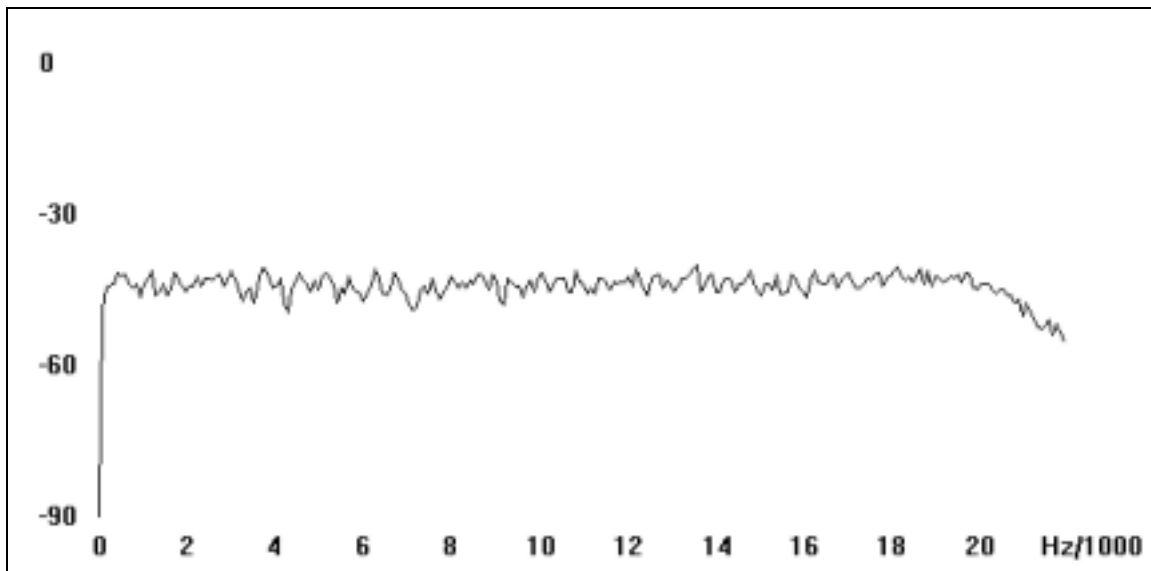


Figure 6 Spectrogram Plot of 20 KHz Noise from Broadband Noise Generator

(Since noise is a statistical process and changes from instant to instant, the “jagged” fluctuations of several dB seen in Figure 6 are not relevant.)

Whenever we discuss a signal to noise ratio, we must specify the noise bandwidth. It would be possible to use the full 20 KHz noise bandwidth, but this results in results that may mislead more than illuminate. Few of us use receivers with a 20 KHz bandwidth for MFSK16 operation. It would be better—but still ultimately misleading—to use a typical receiver bandwidth, say 2.4 KHz. We obtain more information if we instead look at the noise power in the minimum necessary bandwidth to receive MFSK16, which is stated by the MFSK16 developers as 316 Hz. By using this reference, we can look at system

performance using the absolute minimum possible bandwidth. Moreover, should we need to consider other bandwidths, a simple adjustment can be made.

Since the Gaussian noise power is uniformly distributed over the 20 KHz band, the noise power in a smaller bandwidth is given by the ratio of the bandwidths. For example, any 1 KHz segment of the band 0...20 KHz has $1/20^{\text{th}}$ of the energy contained in the entire 20 KHz bandwidth. To convert to dB, simply take $10 \cdot \log_{10}$ of the bandwidth ratio.

Therefore, to adjust the noise power in a 20 KHz bandwidth to that in a 316 Hz bandwidth, an adjustment of $10 \cdot \log_{10}(20000/316) = 18 \text{ dB}$ is necessary. Thus, the noise into the MFSK16 receiving soundcard in a notional bandwidth of 316 Hz, using the parameters shown in Figure 2, is $-10 \text{ dBm} - 18 \text{ dB} = -28 \text{ dBm}$.

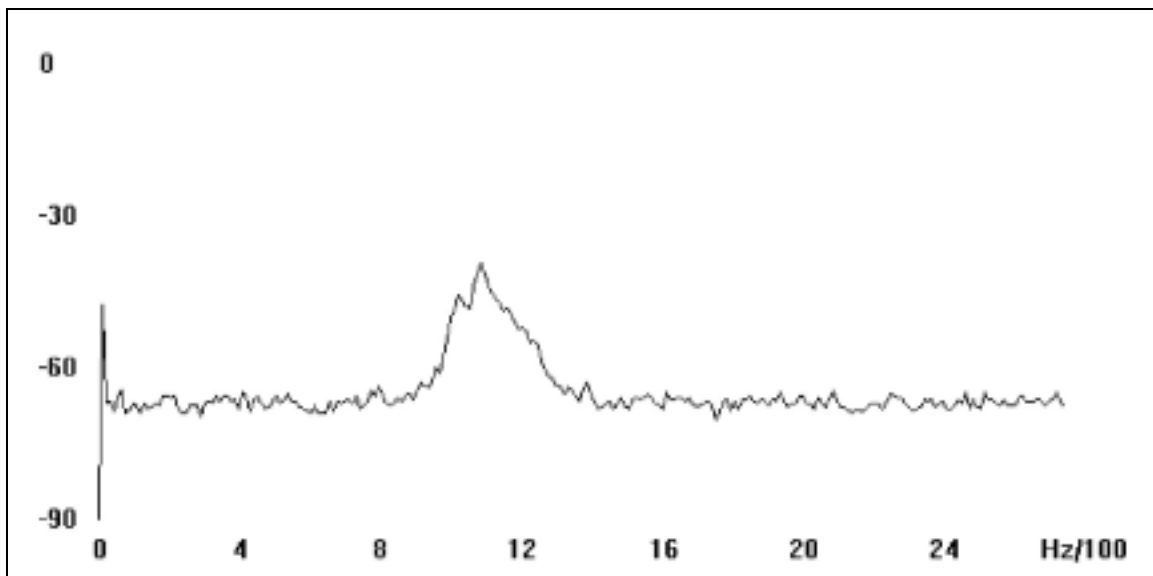


Figure 7 MFSK Signal in Broadband Noise – S/N 10 dB for 20 KHz BW, 2.8 KHz Displayed

Figure 7 shows an expanded spectrum plot of 20 KHz broadband noise at a level of -20 dBm and a MFSK16 signal at -10 dBm , corresponding to a signal-to-noise (S/N) ratio of 10 dB considered over the full 20 KHz noise bandwidth. The MFSK signal stands out because the broadband noise energy is distributed over a full 20 KHz bandwidth, whilst the MFSK16 energy is concentrated into 316 Hz of spectrum. (The Spectrogram program is not calibrated in dBm, so the dB scale indicates relative levels only).

Stream and HamScope performed essentially equally when exposed to Gaussian noise.

GAUSSIAN NOISE PERFORMANCE

SOFTWARE	PRINTING THRESHOLD 316 HZ BW
Stream v. 0.88a	-6 dB
HamScope v. 1.2.1	-7 dB

The negative value in the above table means that the MFSK signal was *weaker* than the noise power by the stated value. To produce good print when the signal is weaker than the noise power in the required bandwidth is impressive performance.

In both Stream and HamScope, only two dB separated excellent print from unusable print, as discussed in earlier.

As a double check on the measured performance, it is worthwhile to examine the S/N ratio for the individual tone filters. The bandwidth of the 16 tone filters is not specified in the MFSK16 documentation, but we know the tone spacing is 15.625 Hz. If we assume a bandwidth of 80% of the channel spacing—a typical value for multi-channel decoding—the individual tone filters have a 12.5 Hz bandwidth. In this case, the noise power in each tone filter is $10 \cdot \log_{10}(20000/12.5) = 32$ dB below the full 20 KHz noise power, instead of the 18 dB correction for 316 Hz bandwidth derived above. Thus, another 14 dB of noise power adjustment should be made when considering the individual tone filters. The signal to noise ratio for the individual tone filters in Stream is, therefore, $-6 + 14 = +8$ dB. This is a reasonable value for proper bit decoding.

Single Tone Interference

The penultimate test run involved a single tone interferer, injected using the test configuration shown at Figure 3.

The frequency of the interferer was adjusted between 700 Hz to 2000 Hz. (The MFSK16 signal had an idle tone of 1000 Hz). Interfering test signals were also injected using frequencies within the MFSK16 modulation envelope itself. A frequency step of 100 Hz was used, except near the MFSK16 occupied bandwidth, where a 50 Hz step was used.

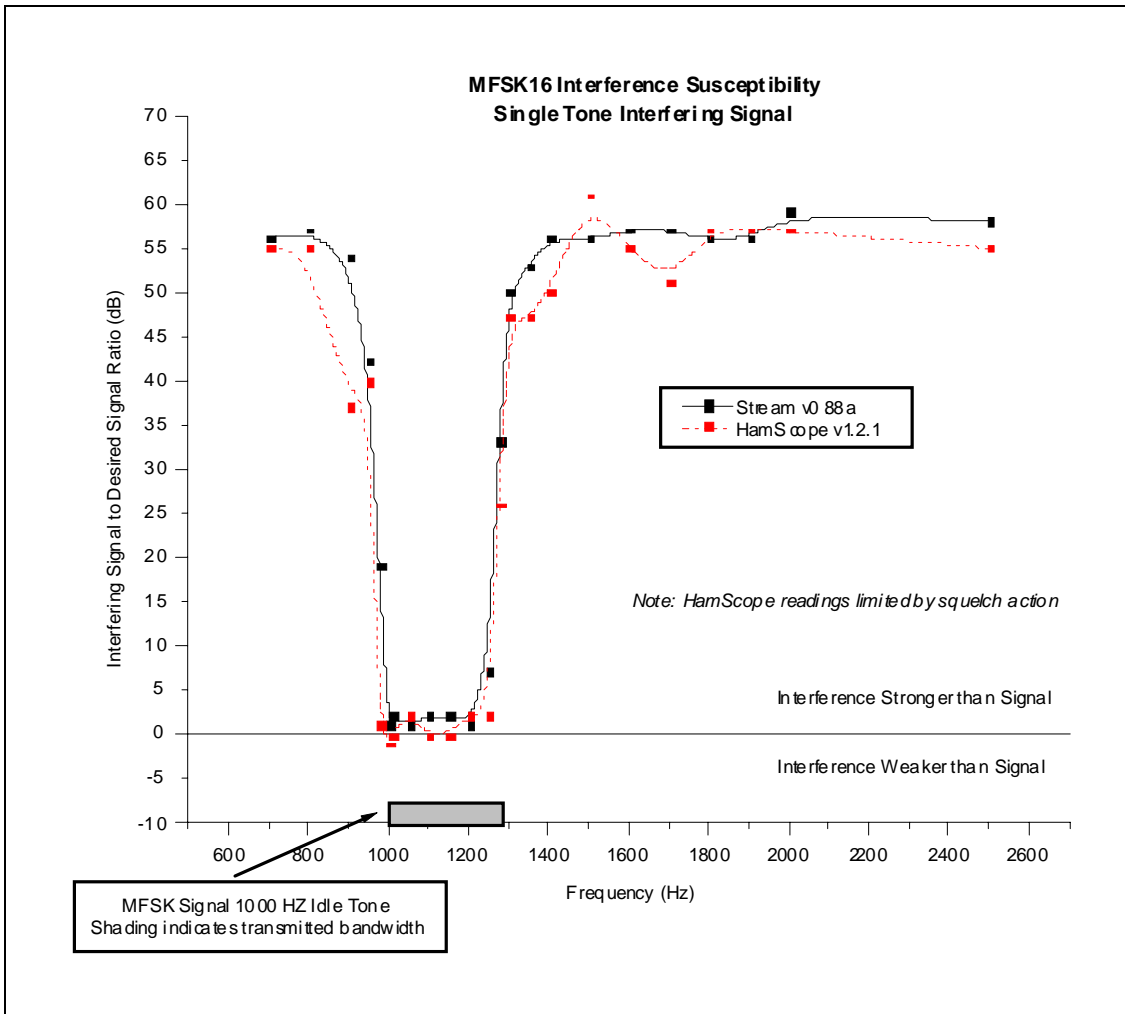


Figure 8 Interference Rejection for Single Tone Interferer

Figure 8 shows the desired-to-undesired signal ratio for no-print threshold for a single tone interferer with varying frequency. For example, an interfering signal at 1700 Hz must be nearly 55 dB stronger than the MFSK16 signal. In all cases, the MFSK16 signal used a 1000 Hz idle frequency. The shaded rectangle indicates the approximate occupied bandwidth of the MFSK16 signal.

At more than 100 Hz offset, the required C/I ratio stabilizes at about 50 to 55 dB. One problem noted with HamScope was the squelch control seemed to be “grabbed” by the interfering tone signal and stopped print, even though the squelch slider was set at minimum. Temporarily removing the interfering tone allowed HamScope to continue printing, thus indicating that lack of adequate desired signal level was not causing the squelch lockup.⁸

⁸ Stream was run with squelch disabled.

Interestingly, even where the interfering signal is *within* the MFSK16 occupied bandwidth, a single tone roughly equal—or even 1 or 2 dB *stronger* than the MFSK16 signal—does not bother print. With a less robust mode, one normally finds that where the interfering signal is within the passband, the interfering signal must be weaker than the desired signal.

The 50 to 55 dB ultimate rejection of both MFSK16 programs is puzzling. A value much closer to the dynamic range of the software, 85 dB or so, would be more reasonable. Without more knowledge of the “internals” of Stream and HamScope, it isn’t possible to definitively ascribe reasons to the difference between the dynamic range and the ultimate rejection, but the author suspects soundcard limitations, not software design issues.

Since the data shows that the interference potential of the single tone is relatively constant once it is more than a hundred Hz or so from the MFSK signal, a single tone interferer is a useful model for CW or PSK31 or other narrowband modulation modes. In other words, it is the interfering energy, not necessarily its precise frequency, that is the problem.

MFSK16-to-MFSK16 Interference

The final series of tests used an MFSK16 signal as the interferer. Here, we have a greater potential for partially garbled, but still useful, print. The interfering signal spreads its energy over all 16 tones, as, of course, does the desired MFSK16 signal. At any instant only one tone of the interfering signal will be active, thus statistically distributing the interference. Unless the interfering MFSK signal is exactly aligned in frequency with the desired MFSK16 signal, a portion of the interfering signal will therefore fall outside the passband of the desired signal. In addition, the interfering MFSK16 signal will be asynchronous with respect to the desired signal, thereby providing additional interference rejection for the integrate and dump receiving filters.

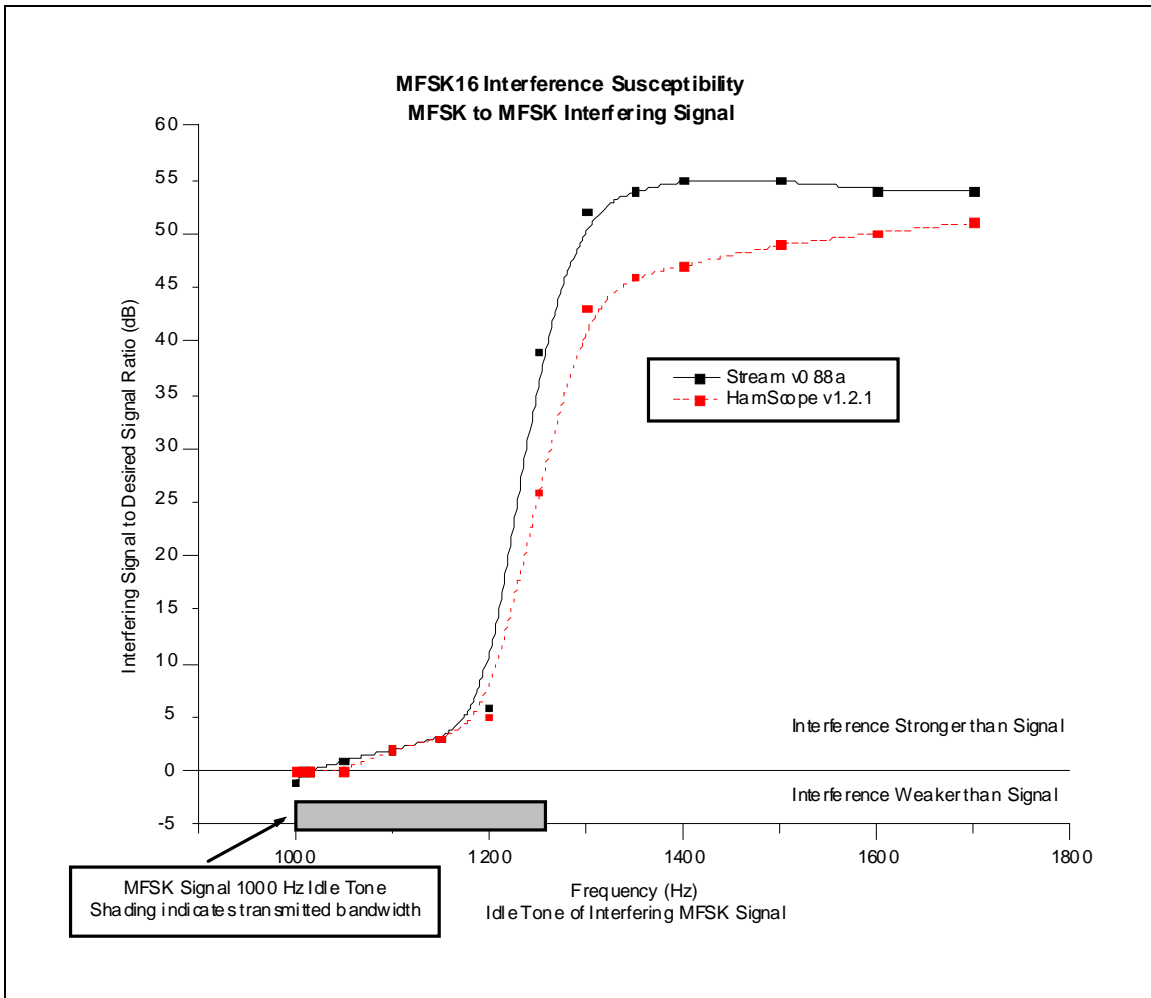


Figure 9 MFSK-to-MFSK Interference Rejection

Where the desired and interfering signals have very small offsets (between 0 and 200 Hz), the energy of the desired and interfering signals actually overlaps. (Since the interfering MFSK16 signal is modulated, part of its energy fall outside the passband of the desired signal unless they are exactly aligned in frequency.) Even in this case, however, MFSK16 produces excellent print, so long as the desired and interfering signals are about equal in strength. Once the interfering MFSK16 signal is offset sufficiently to remove this energy overlap, the required S/N stabilizes at about 50 to 55 dB, about the same as for a single tone interferer. This, of course, is as expected, since once the energy overlap of the two MFSK16 signals are disentangled, the interfering MFSK16 signal just looks like a series of single tone interferes.

Here again, the HamScope software did not perform as well as Stream, although it did not exhibit the squelch lockup observed in the single tone interference study.

So, What Can I use Immediately Out of All this Stuff?

Assuming that your soundcard performs about as well as the one studied and if you are using a SSB receiver set for voice reception, a nearby signal 55 dB stronger than the desired signal, even though it is several hundred Hertz away from the desired MFSK16 signal, will cause loss of print.

In the best circumstances, better receiver selectivity would be used to knock down the interfering signal. Unfortunately, many receivers permit using only an SSB bandwidth filter when receiving MFSK16. Substantially better receiver performance will be found if a 400 to 500 Hz wide IF filter is used for MFSK reception.

There are two inter-related potential interference problems affecting MFSK16, AGC capture and filter rejection. Although each is discussed separately, both events occur simultaneously when a desired and an interfering signal are present within the receiver passband.

AGC Capture

AGC action is very complex, if analyzed in detail. The following discussion takes some liberties with the precise details of AGC operation, but preserves the essential effects of widely differing signal levels.

Recall that AGC works by reducing or increasing the gain of the receiver to maintain a constant audio output as the input signal varies. If two signals are simultaneously in the passband of an HF receiver, the automatic gain control (AGC) will attempt to keep the strongest signal at a fixed level by reducing gain.⁹ Of course, a receiver can't simultaneously reduce gain for one signal and increase gain for a second. Consequently, the recovered audio for the weaker signal will be reduced. At some point, the desired signal gets lost in the residual receiver noise. The reduction in gain of the weaker signal equals the difference in signal level between the two signals.

A simple example will illustrate AGC capture. Suppose you are receiving one signal at a received signal level of -90 dBm, and that the resultant audio output level for this signal is -10 dBm. A second signal comes on, within the passband of your SSB receiver. Not so close as to cause interference, but still within the receiver passband. The second signal is stronger, say 30 dB stronger at -60 dBm. In the absence of AGC action, the second signal would be +20 dBm audio output. But, the receiver's AGC will apply 30 dB of gain reduction so that the second signal produces an audio output to -10 dBm. To maintain this audio level with a 30 dB stronger signal requires 30 dB gain reduction. This 30 dB reduction, of course, also applies to the desired first signal, so its audio output is reduced from -10 dBm to -40 dBm.

We know that both Stream and HamScope software running on a TI5200 computer have a dynamic range is 85 dB. Therefore, the AGC gain reduction—30 dB in the example—

⁹ Actually, it keeps the sum of the signals constant. However, in most cases one signal will be dominant and will govern AGC action.

will not cause any reduction in print quality. The -40 dBm audio signal still produces perfect print. If the undesired signal is more than 85 dB stronger than the desired signal, however, the AGC capture will drop the desired signal below the dynamic range threshold of either Stream or HamScope software, even if an audio filter is in use.

Moreover, the 85 dB dynamic range measured is based upon a CD-ROM player, a device optimized for high fidelity audio applications with very low noise, and a high quality 600-ohm attenuator. The noise floor performance of a typical SSB receiver can be expected to be significantly inferior.

Software Filter Performance

When examining interference rejection using a single tone interferer and the methodology illustrated at Figure 3, receiver AGC considerations are irrelevant—the desired signal remains well above the dynamic range floor of the test software, even if it is 60 dB below the interfering signal. Here, the interference effect is either inadequate filter rejection in the software filters implemented, or possibly DSP-based effects in the receiving soundcard.

Combined Effect

The combined effect of AGC capture and software filter performance is to simultaneously reduce the strength of the desired signal and to expose it to a much stronger interfering signal. The results presented in this section reflect both effects, at least as long as the interfering signal is within the receiver passband.

In the HF1000 and the Kenwood 870 narrowband tests, in-passband effects are not seen for interfering signals far removed from the receiver bandwidth. Further, the interfering signal strength is reduced by the selectivity of the receiver filters. Here, the print problems are likely due to receiver effects.

Proper Receiver Bandwidth Selection

In considering interference rejection, it is worthwhile to again reiterate the importance of using a proper IF filter in the receiver. Much greater interference rejection—not to mention increased operator enjoyment—will be found if equipment is used that permits a 400 to 500 Hz IF filter for MFSK16 reception.¹⁰ Indeed, with a good IF filter it should be possible to operate within a few hundred Hz of an MFSK16 signal at 80 dB or more signal level difference. Use of a standard SSB bandwidth receiver must be regarded as an unfortunate necessity in many cases. Further, as shown below, adding an auxiliary audio filter does not materially improve interference rejection.

Finally, these measurements do not show HamScope to be materially worse than Stream. Except in extreme cases, both performed equally. The user's selection of HamScope or

¹⁰ The author's Kenwood 870, for example, permits independently adjusting the lower and upper passband limits for SSB. For MFSK16 operation, normally the lower pass limit is set to 1.000 KHz and the upper pass limit is set to 1.400 KHz, yielding a 400 Hz passband. However, the filter skirts in this configuration are not as sharp as found in the same receiver in CW or RTTY Mode. See later discussion.

Stream should be made on other grounds, such as layout, flexibility or ease of use concerns.

RF Verification Measurements

To verify the accuracy of the foregoing discussion of AGC interaction with interfering signal, a separate series of tests were conducted using the setup shown in Figure 10.

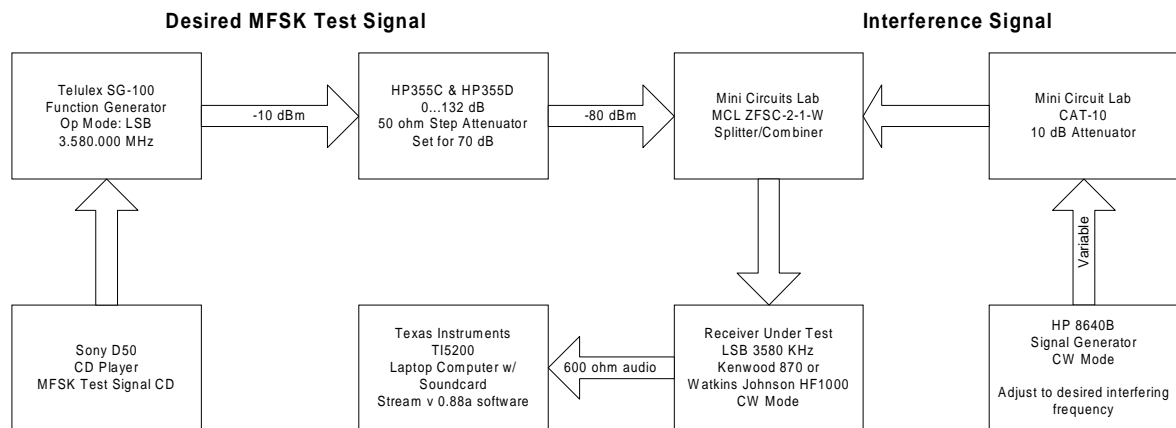


Figure 10 RF Measurement Test Setup

This test setup duplicates receiving an MFSK16 signal off the air, with a potential steady tone (CW) interferer.¹¹ Using this configuration, the frequency and output level of the HP 8640B signal generator was varied to simulate a single tone interfering signal. The desired MFSK signal was generated using the test signal CD described earlier. The particular MFSK transmit frequency selected was 3580 KHz, LSB mode. The MFSK test signal used a 1 KHz idle tone, corresponding to a RF frequency of 3579 KHz. The interfering signal generator was set at a series of frequencies yielding an interfering beat frequency of 500 to 2000 Hz, in 100 Hz steps.

For each interfering frequency, the interfering signal was reduced to a level where perfect MFSK16 print was achieved, then gradually increased until intelligible print ceased. The ratio between the desired signal (held constant at -80 dBm) and the interfering signal can then be simply determined by subtracting the HP 8640B signal generator output, considering the 10 dB fixed attenuator, from the desired signal.

Some data was taken with a Panasonic VP-8191A synthesized signal generator instead of the HP 8640B. It was found that sideband noise in the VP-8191 limited rejection

¹¹ Review of the data shows that once the interfering signal energy does not overlap the MFSK16 spectrum, the interference potential of the undesired signal is relatively independent of the precise offset frequency. Thus, an equivalent power continuous tone interferer can accurately model a modulated signal, such as PSK31, that spreads its energy over a range of frequencies, where the interfering signal does not actually overlap the MFSK16 spectrum.

measurements to about 65 dB. Accordingly, the narrow bandwidth data, where rejection was expected over this level were taken with the HP8640B. The output levels of the various signal generators were set using a HP 3400A true RMS meter. A 50-ohm through termination was used at the HP 3400A. As with the earlier tests, since the same meter was used to set all generators, only attenuator error and splitter error will affect the signal ratio. Earlier checks of the attenuators and splitter with an HP 8754A network analyzer established the accuracy of the attenuators and splitter as within 0.5 dB.

This series of tests were repeated for 5 receive configurations:

1. Kenwood 870 transceiver, LSB mode, low frequency cut = 1000 Hz, high frequency cut = 1400 Hz (400 Hz bandpass), no audio filter.
2. Kenwood 870 transceiver, LSB mode, low frequency cut = 300 Hz, high frequency cut = 2800 Hz (2500 Hz bandpass), with a Datong FL3 outboard audio filter, set to pass a 400 Hz bandwidth, centered around the MFSK16 received tones.
3. Kenwood 870 transceiver, LSB mode, low frequency cut = 1000 Hz, high frequency cut = 1400 Hz (400 Hz bandpass), with a Datong FL3 outboard audio filter, set to pass a 400 Hz bandwidth, centered around the MFSK16 received tones.
4. Kenwood 870 transceiver, LSB mode, low frequency cut = 300 Hz, high frequency cut = 2800 Hz (2500 Hz bandpass), no audio filter.
5. A Watkins Johnson HF1000 receiver, CW mode, bandwidth 325 Hz, BFO offset to center MFSK16 tones in the passband. No audio filter.

Configuration 4 should track, within the accuracy of the test equipment and protocol, the results obtained with the setup of Figure 3, since the receiver selectivity does not reduce the interfering signal. Differences between Configuration 4 and the other receiver arrangements will, of course, reflect differences in receive selectivity and receiver architecture.

Figure 11 shows the interference rejection of each Configuration.

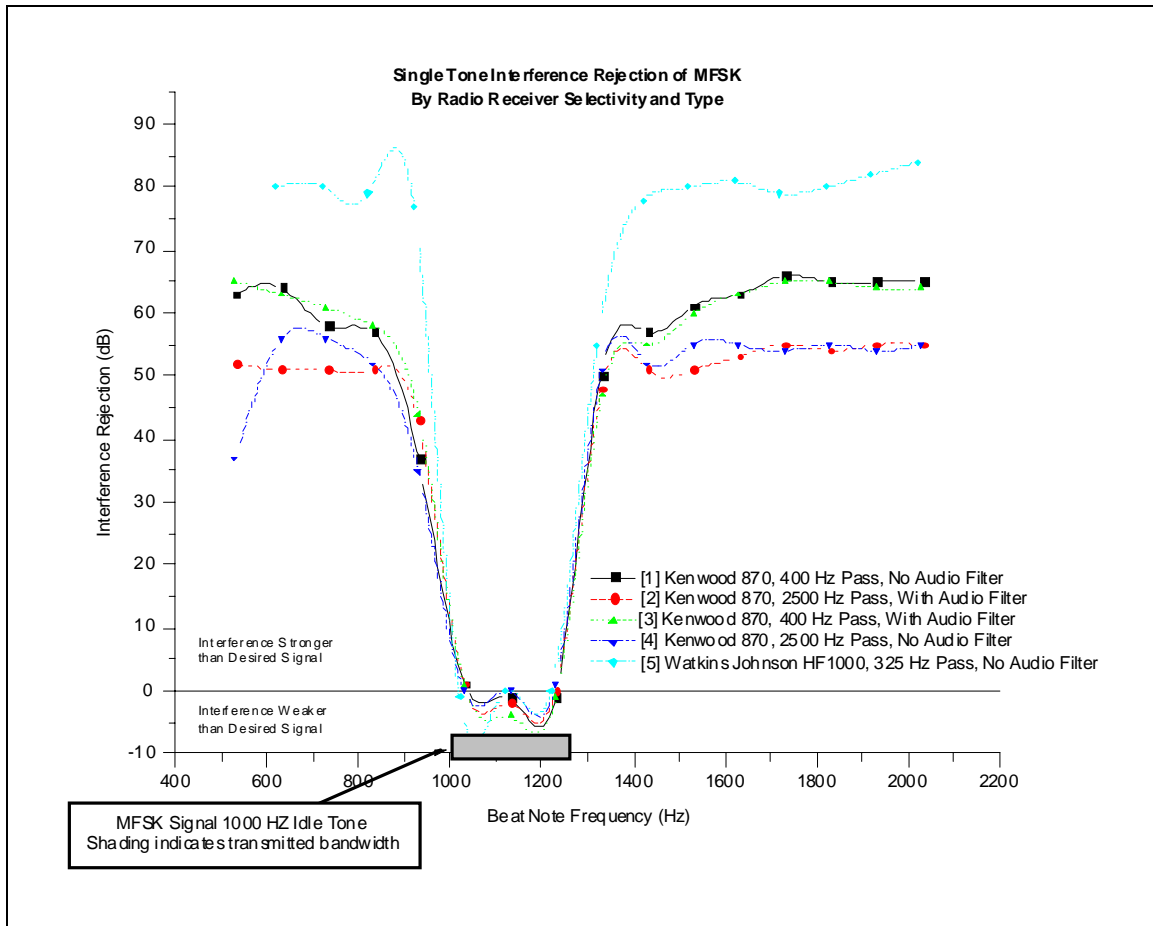


Figure 11 – Single Tone Interferer Performance, Effect of Receiver and Audio Filter Selectivity

Studying Figure 11 teaches several lessons:

- An anomalous data point is seen for about 550 Hz beat note, case [4]. This is due to the second harmonic of the beat note falling within the MFSK16 passband. Thus, the observed interference rejection of about 37 dB is really a measure of the second harmonic distortion in the Kenwood 870 receiver.
- When the interfering tone falls within the MFSK16 passband, receiver selectivity doesn't matter. All receivers perform about the same, requiring the interference to be a few dB weaker than the desired signal.
- The performance of the software filter incorporated in the Stream software governs interference performance when the interfering tone is within the MFSK16 passband or very close to the passband. Thus, all 5 plots are essentially identical in the range out to about 50 Hz on either side of the required MFSK16 passband.
- Using an outboard audio filter doesn't help interference rejection as much as expected, a counter intuitive result. It does clean up the second harmonic problem for interfering tones in the 500-600 Hz range, and it does slightly improve interference rejection for tones in the 900 – 1000 Hz range (for the desired

MFSK16 signal with a 1000 Hz idle tone). This suggests that the strong interfering signals are causing AGC gain reduction—or other—problems based upon the roofing filter, not the Digital Signal Processing (DSP) filter in the 870.¹²

- As is normally the case in filters, rejection on the high side is better than on the low side.
- Comparing Figure 8 with the Kenwood SSB wide data (curve 4) of Figure 11 show about the same interference performance, as expected.
- The Watkins Johnson HF1000 is a much better performer than the Kenwood 870, even though set for similar bandpasses (325 Hz vice 400 Hz), showing more than 20 dB better interference rejection.

Figures 12 and 13 illustrate the effective bandpass differences of the Watkins Johnson at 325 Hz (Figure 12) and the Kenwood 870 set for 400 Hz bandpass (Figure 13). The figures are not true bandpass sweeps, but are instead use Spectrograms to plot the received noise level. This technique provides a quick illustration of the passband performance of a receiver, without the time required for a point-by-point frequency sweep. Although it fails to show the ultimate rejection, the differences in the relative selectivity of the two receivers is clearly shown. The Watkins Johnson has a much superior bandpass, with excellent skirt selectivity. The Kenwood 870, in contrast, has much gentler skirts. This is likely a deliberate choice by Kenwood, as sharp cutoff filters may produce less pleasing sounding SSB. Examination of the RTTY or CW modes in the Kenwood 870 shows a filter that more closely resembles the Watkins Johnson. Figure 14 shows the Kenwood 870 in 500 Hz RTTY mode. It is to be regretted that Kenwood has not chosen to allow a more flexible selection of receive filters in the 870 whilst transmitting in SSB mode.

¹² Both the Kenwood 870 and WJ HF1000 use IF DSP. However, the HF1000 has much more flexibility on filter bandwidth and BFO offset.

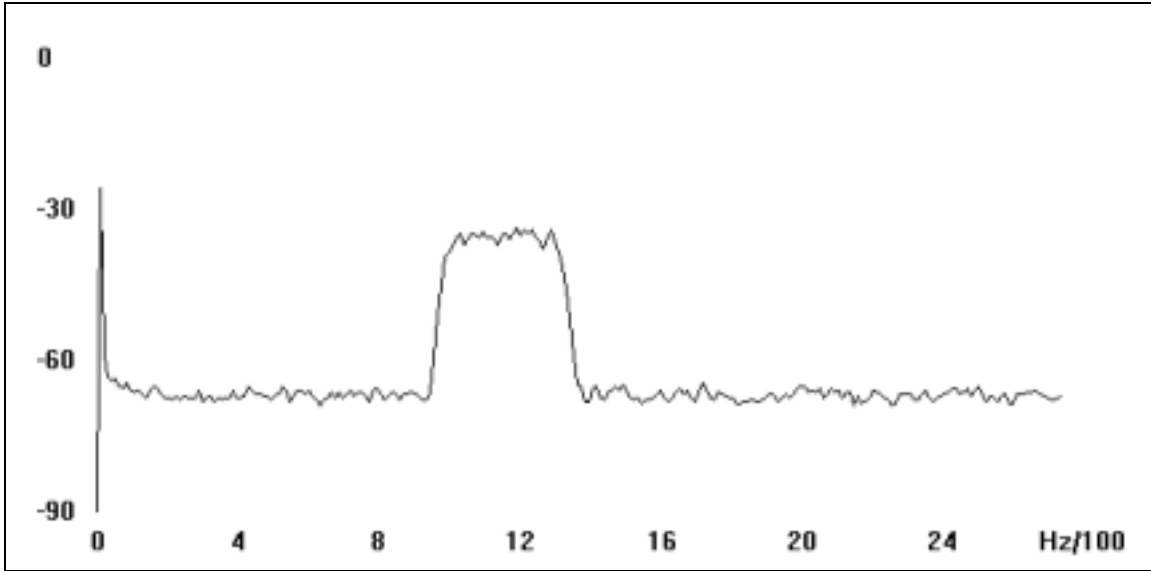


Figure 12 Watkins Johnson HF1000 Receiver, Bandwidth = 325 Hz Noise Plot

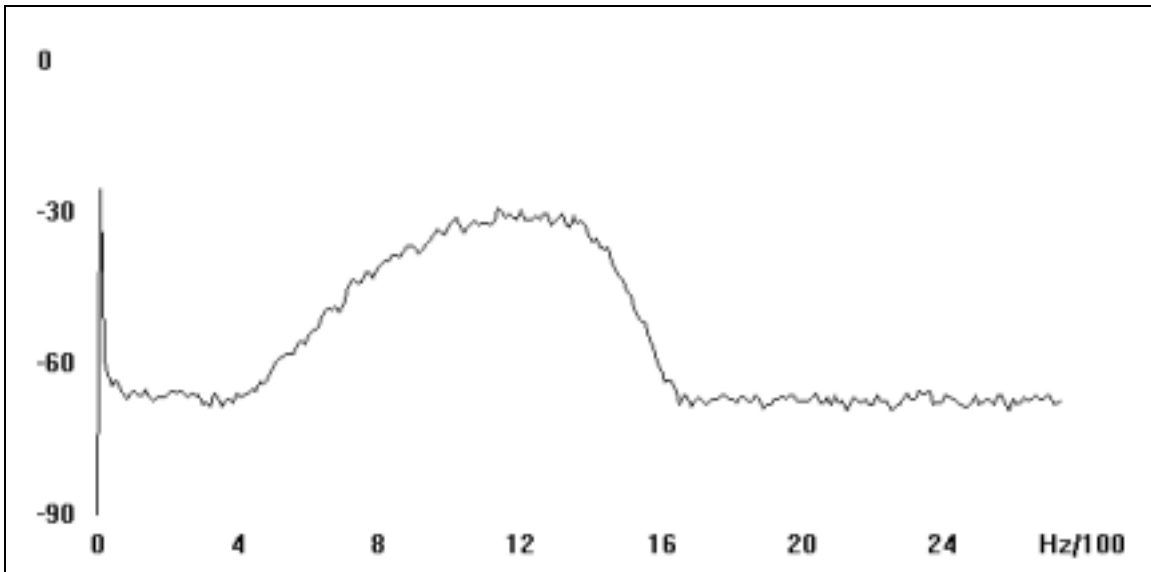


Figure 13 Kenwood 870 Receiver 1000 Hz Low Cut 1400 Hz High Cut SSB Mode Noise Plot

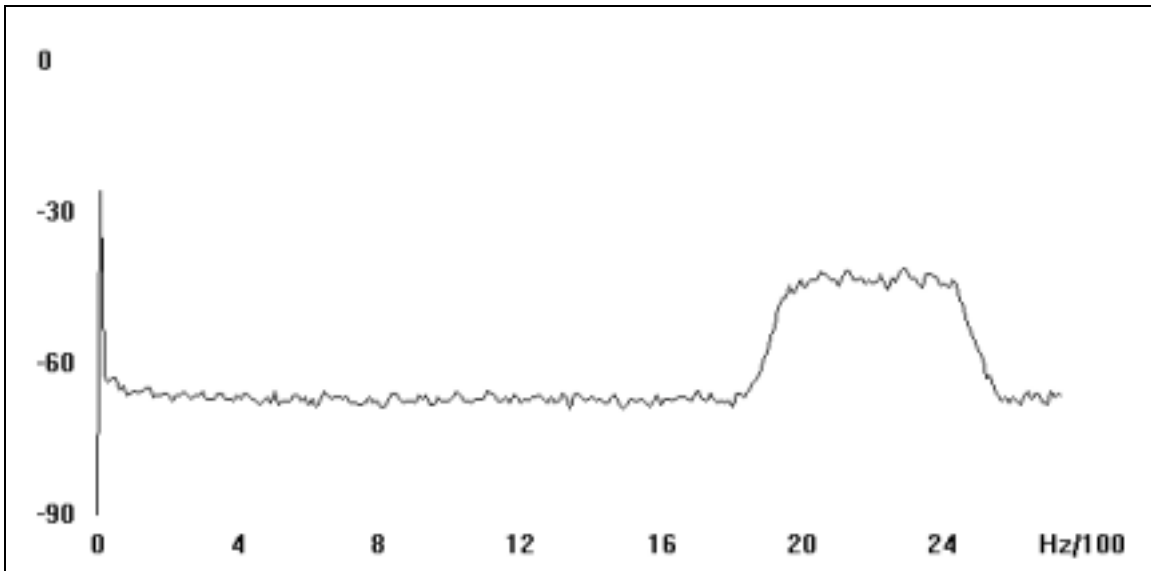


Figure 14 Kenwood 870 FSK Mode, 500 Hz Bandwidth High Tones Noise Plot