LO Phase-Noise Management in Amateur Receiver Systems

Rick Campbell KK7B
TriQuint Semiconductor
2300 NE Brookwood Parkway
Hillsboro, OR 97124
rickc@TQS.com

When experimenting with microwave radio, there are few pieces of commercial equipment that may be used without modification. Some of us choose to design and build complete stations from scratch, a full time avocation in itself. Others purchase a commercial HF or VHF transceiver and build frequency converters. Even those who purchase commercial transverters often discover that performance can be considerably improved by understanding the receiver system and making a few modifications. This paper presents a practical discussion of microwave weak-signal receiver design, with a focus on frequency conversion, gain and selectivity distributions, and LO phase noise effects.

Multiple Signals, RF Filter Noise Bandwidth, and LO with Phase Noise

The RF, LO, and IF spectra in a typical amateur frequency converter are shown in figure 1. The RF input to the mixer passes through a flat passband RF filter with a few signals above the noise floor. The LO has appreciable phase noise. Convolving (see appendix) the RF input signal spectrum with the LO spectrum results in the IF spectrum shown. The image noise and LO phase noise effects are all well below the receiver noise floor, and will not affect the receiver signal-to-noise ratio for any of the signals in the passband. The desired narrow IF tuning range is also indicated.

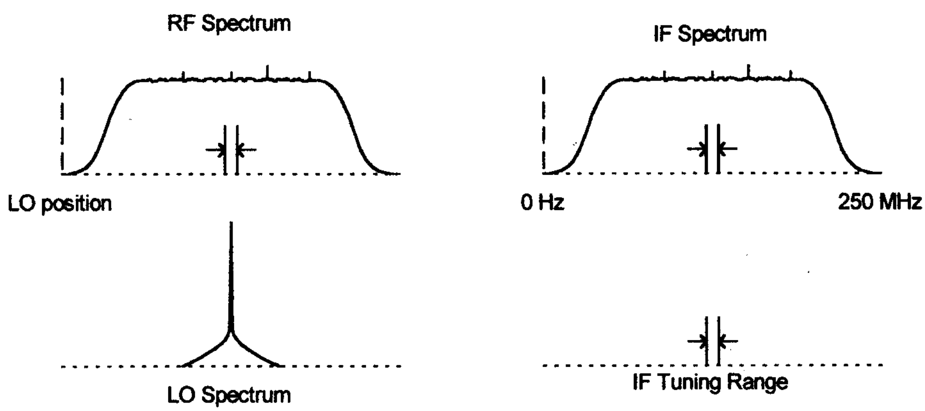

Figure 1

Figure 2 shows the same converter with a strong signal in the RF passband. Note how the LO phase noise modulated onto the strong signal in the IF output spectrum buries the weaker signals in the noise. This is called reciprocal mixing.

From the above discussion, it might be concluded that LO phase noise is only important when trying to hear weak signals in the presence of strong signals. This is generally true for receivers with only one LO. Multiple conversion superhet receivers have more than one LO, and multiple LOs can combine to produce internally generated receiver spurious responses. Proper receiver frequency planning will ensure that these spurs are outside the normal tuning range of the receiver. If they are strong, however, they can still convert LO phase noise into the desired passband and raise the receiver noise floor.

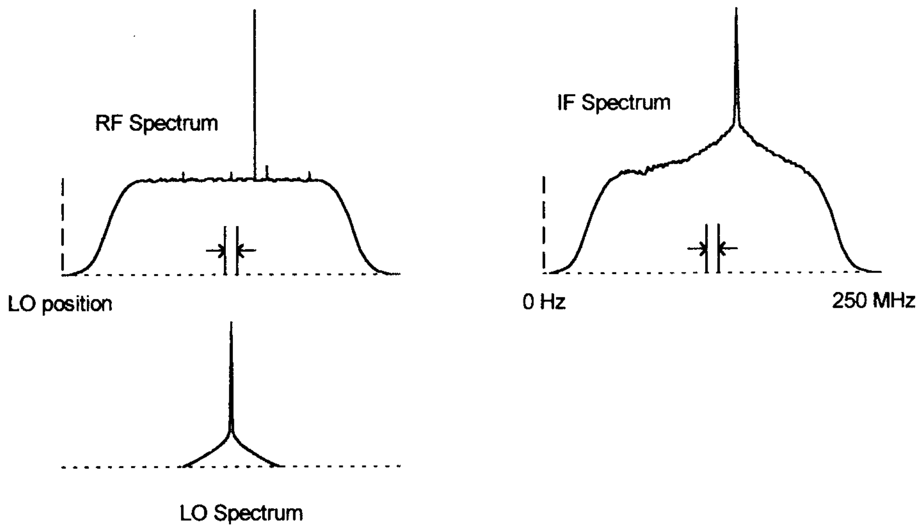

Figure 2

One way to reduce the effects of reciprocal mixing is to minimize the RF bandwidth, so that strong out-of-band signals and spurs are attenuated before the mixer. With a good frequency plan that moves internally generated spurs far from the desired frequency band, and proper RF and IF bandpass filters, even a receiver with modest LO phase noise can have excellent performance. Unfortunately, the trend in HF and VHF transceivers is to offer complete coverage of the widest RF bandwidth possible, which tends to exacerbate phase-noise problems.

Figure 3 illustrates the 2nd conversion using a radio with an RF input filter much wider than the indicated desired tuning range. As long as the LO has low phase noise, there are no problems, and the spectrum at the output of the 2nd mixer in figure three is a faithful replica of the RF input signal.

Figure 4 shows the same filter arrangement in a receiver with poor LO phase noise. Even though the strong signal (or receiver spur) is well outside the desired tuning range, the LO phase noise sidebands significantly increase the receiver noise floor.

Cleaning up the LO phase noise would solve this particular problem, but reducing the bandwidth at the input to the mixer will also eliminate LO phase noise contributions from out-of-band signals, as well as reducing receiver spurious responses and improving the performance of following mixers and amplifier stages. Figure 5 shows spectra in the same receiver, but with a good filter added at the RF input port. The strong out-of-band signal is now attenuated to the point where LO phase noise contributions in the desired tuning range are well below the noise floor.

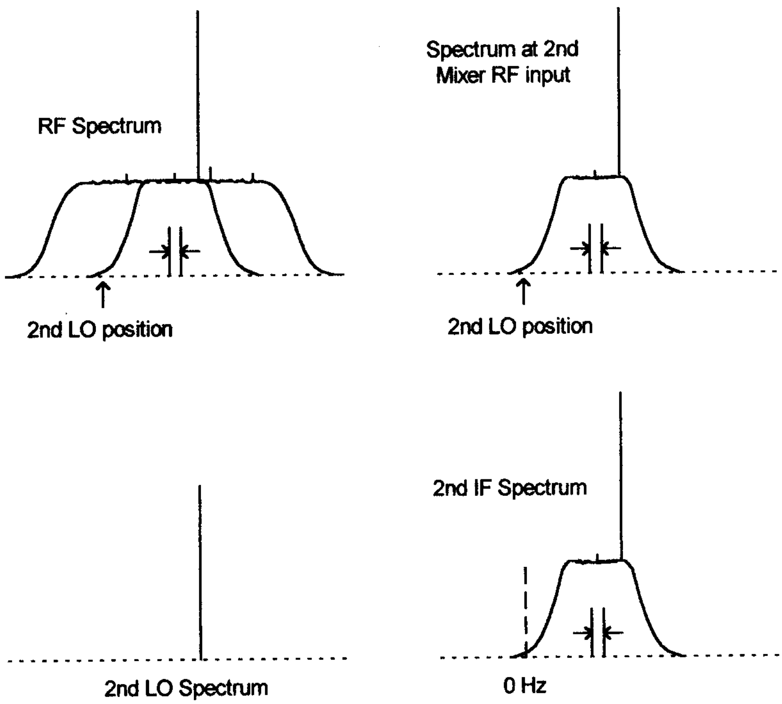

Wide Bandwidth IF Receiver With Clean LO
Figure 3

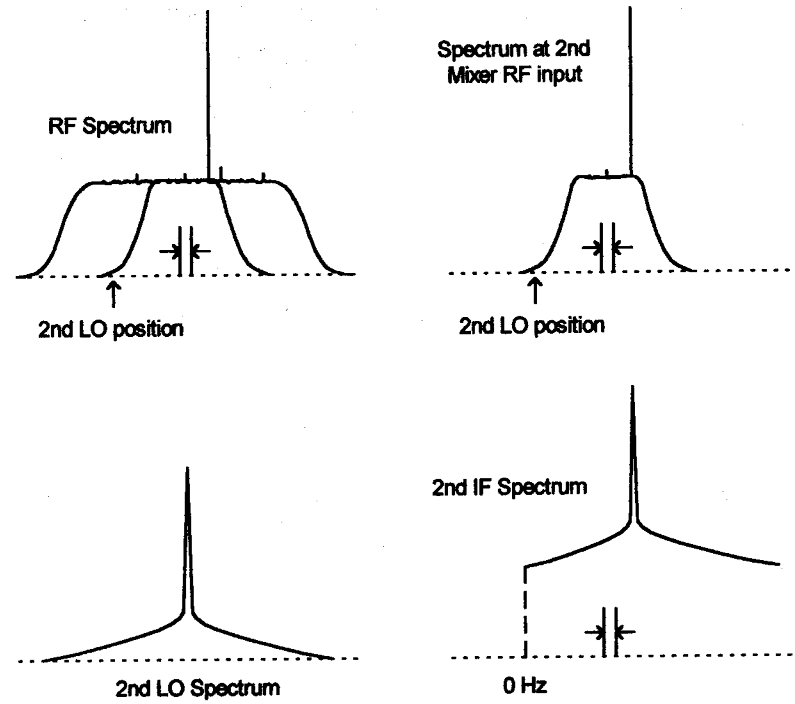

Figure 4

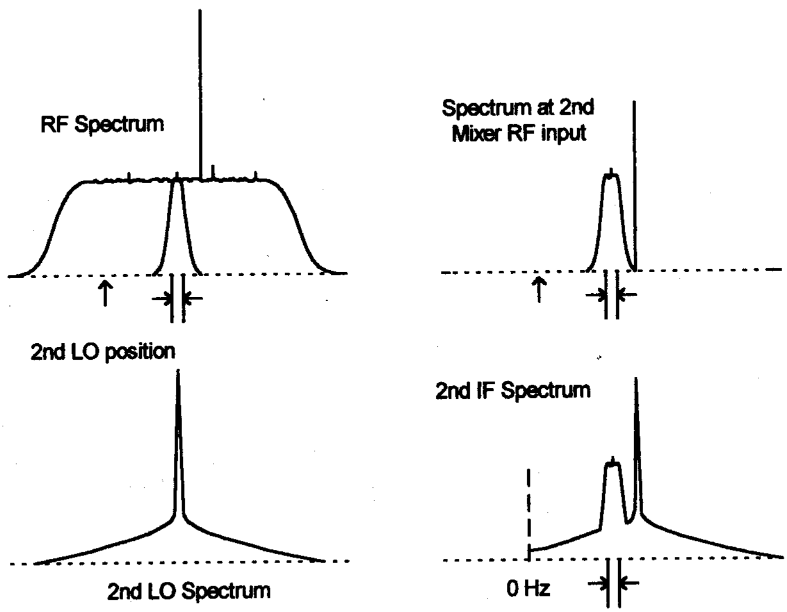

Figure 5

Case Studies

Figure 6 is a block diagram of a 1296 MHz receive converter. The 9th overtone 144 MHz crystal oscillator is multiplied by 8 to obtain the 1152 MHz LO signal. The desired signal at 1296.100 MHz is then downconverted to 144.100 MHz. There is a strong receive spur at 144.000 MHz, where the crystal oscillator signal leaks out of the transverter case and is picked up by the IF rig. A typical crystal oscillator has +7 dBm output, and good shielding would reduce this by over 60 dB at the IF rig antenna port. Suppose the144.000 MHz spur is –60 dBm, and the receiver noise floor in a SSB bandwidth is around –136 dBm. The spur is then 96 dB above the receiver noise floor. A good, wide dynamic range 144 MHz transceiver should be able to easily handle this signal while listening at 144.100 MHz. Most of us would prefer to use the 2m multi-mode transceiver we traded for an aria at the last hamfest.

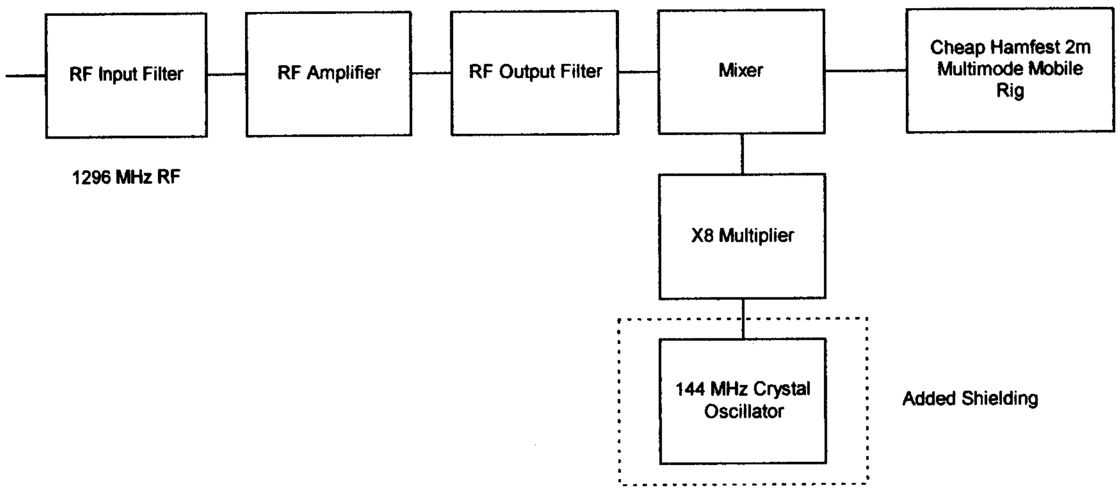

Figure 6

There are three changes that may be made to improve the figure 6 receiver system sensitivity. The easiest is to increase the gain at the front end until a 3 dB increase in the IF rig noise floor is unimportant. This compromises dynamic range, but that may be entirely acceptable in some locations. The second technique is to reduce the level of the 144 MHz birdie at the IF rig. This might be achieved simply by moving the two boxes apart, or by adding some shielding or bypassing inside the transverter case. Multiplying by 6 or 10 instead of 8, so that a 144 MHz spur is avoided entirely, is good practice. Finally, improvements to the IF rig LO phase noise could be made, but improving LO phase noise in a commercial transceiver is a serious project that should usually be considered a last resort.

The system shown is figure 6 has a fairly obvious flaw, and a birdie that is easy to trace. The "phase-noise sensitivity problem" with the system in figure 7 is more insidious. This system has the common 3456 MHz frequency scheme of a 2160 MHz LO for the first conversion to 1296 and an 1152 MHz LO for conversion to 144 MHz. The third harmonic of 1152 MHz is a strong spur right on 3456 MHz. To move it out of band, the 1152 MHz LO frequency was lowered 50 kHz, which put the spur at 3455.850 MHz and moved the IF up 50 kHz. This moved it out of the tuning range on all rigs used as IFs. The system worked fine with a low phase-noise ICOM 202 as an IF, but had a very high noise floor with a synthetic radio. The strong 3455.850 MHz signal is generated in the 2nd mixer, and was only attenuated about 30 dB by the 1296 MHz IF amplifier and filter. Putting the 1296 MHz to 144 MHz converter in a shield box and adding an interdigital filter in the 1296 MHz IF line completely cured the problem.

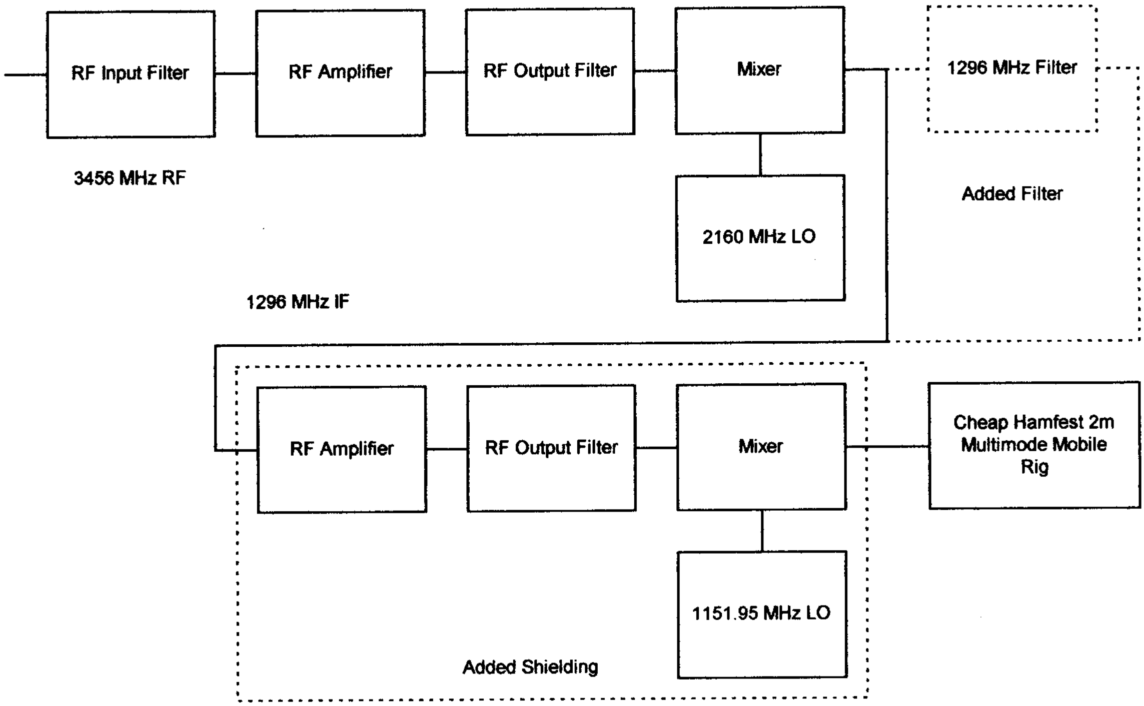

Figure 7

Note that in both of the above situations, reduced sensitivity due to "poor IF rig LO phase noise" could be cured by understanding what was happening and making simple modifications to the rest of the system.

Good Practice

Here are a few good practices and general guidelines that can help avoid LO phase noise problems in amateur microwave systems.

1. Avoid LOs with high phase-noise. This one is obvious, but the fact is that the first LO in a microwave system has the biggest impact, and it is usually one that we have some control over.
2. Use narrow bandwidth filters attached to the RF port of every mixer in the system. This will prevent signals well outside the frequency range of interest from raising the receiver noise floor by reciprocal mixing with down-stream LOs. It is very good practice to always include a narrow filter in the IF transmission line to a 2m or HF transceiver. Excellent models are available commercially from XXXX, surplus filters

from mobile radios can be retuned, and homebrew three resonator filters can be built using the article by Hayward [2].

3. Avoid frequency plans that produce strong close-in birdies. These include most of the clever multi-band microwave schemes published by hams over the decades.
4. Minimize the number of frequency conversions.
5. Consider experimenting with a Direct Conversion receiver. They do not generate internal spurs.

Appendix: A Frequency Conversion Tutorial

Rick Campbell
9/9/99

Frequency conversion is a physical process that may be approximately described mathematically. The description can include equations and graphs, two of many mathematical dialects. In this paper, we will describe concepts and introduce a graphical technique, but avoid equations. The references include both graphs and equations, and are recommended for anyone inclined to pursue frequency conversion, signal-to-noise ratio, random processes, and phase-noise effects on an analytical level. Such analysis is a full-time avocation, however, and will generally impede efforts to get a signal to the moon and back.

Convolution

We can describe signals in the frequency domain or the time domain. We can examine the same signal on an oscilloscope in the time domain and on a spectrum analyzer in the frequency domain. Engineers like to keep these separate, and use Fourier integrals to transform back and forth between domains. Typical amateur signals, particularly those buried in noise and bounced off moving targets or propagated through cities, are exceptionally difficult to handle using engineering mathematics. Fortunately, we can understand and use the mathematical concepts while avoiding the ponderous details and approximations inherent in the full analytical approach.

A frequency converter (mixer) multiplies two signals together in the time domain. Multiplication in the time domain is equivalent to convolution in the frequency domain. Suppose we look at two different signals on a dual-trace oscilloscope, and then multiply them together and look at the product on another oscilloscope. This is the time-domain view. We could actually pick off individual points on the two original signals, multiply the numbers together, and obtain the product waveform. We could also look at the same two signals on a spectrum analyzer, and then the mixer output on another spectrum analyzer. From the ARRL Handbook, we know that we would see "sum" and "difference" frequencies—but what if one or both of the signals is not a simple sine wave? Convolution gives us a rigorous mathematical way to predict what we will see on the output spectrum analyzer.

Convolution involves sliding two complex frequency domain signals past each other, multiplying and adding. The result of convolution in the frequency domain is a plot of complex amplitude versus frequency. The frequency scale is the offset between the two spectra, and the amplitude scale is related to how well the two signals line up at each offset frequency. In many cases, particularly those involving uncorrelated signals and noise, we can learn much about the signals by sliding their power spectra past each other and multiplying, rather than performing the complex arithmetic. The power spectrum is what a spectrum analyzer displays.

To illustrate frequency sliding, consider two filter passbands in a receiver with passband tuning and a pair of narrow CW filters. Connect an audio voltmeter to the receiver output and disconnect the antenna so there is only white noise in the receiver. Adjust the passband tuning to one extreme, where the two filter passbands do not line up. The receiver noise output will be very low. Now slide one filter passband past the other by slowly turning the passband tuning control through its range. When the passband edge of one filter begins to overlap the passband edge of the other, the noise begins to increase. The noise is at a maximum with the filters are perfectly lined up. If the two filters have identical, rectangular passbands, the plot of noise output versus offset frequency will have a triangular shape, as shown in figure A1. If one filter is wide and one is narrow, the shape will be flat on top, but with triangular sides. If one filter is very narrow and rectangular, the shape will be almost indistinguishable from the shape of the wide filter. Figure A2 illustrates the output that results from several different combinations of filter shapes. If the two passbands are centered at different frequencies, then the frequency scale is shifted, and maximum output occurs at an offset equal to the difference between the two center frequencies.

Of particular interest is the convolution of the spectrum at the output of a filter with a spectrum containing just a single frequency. The output is just the original spectrum, shifted to the sum and difference frequencies. Shifting an amateur band to a new frequency range is exactly what we want to accomplish in our frequency converters. This process is called "heterodyning," and circuitry and simplified theory are discussed at length in every radio handbook. The handbook theory does not provide a simple way to describe what happens when a Local Oscillator has a spectrum that includes noise and spurious outputs. Convolution provides a mathematically rigorous way to describe what happens when an arbitrary LO multiplies a frequency range containing multiple signals and noise.

The mathematical development of graphical convolution is provided in a number of textbooks. The treatment by Ambardar is particularly recommended [1]. We can approximate a graphical convolution by making two tracings of signals from the Spectrum Analyzer screen and sliding them past each other to find out how much output we get at a particular frequency offset. The examples below illustrate frequency conversion using various LOs.

Clean LO

Figure A3 shows two signals 30 dB above the noise floor at the RF input of a mixer, with a clean LO signal at the LO port. After sliding the RF port signals past the LO port signal, we obtain the IF port signal. The clean LO picks out the RF signals at the appropriate frequency offsets, and picks out noise at all other frequencies.

LO with Discrete Sidebands

Figure A4 is the same situation, but now with a dirty LO. The clean LO signal still picks out the two desired signals, but now the LO spurs also pick out signals. The IF output has the two desired signals, and weaker desired signals offset by the LO sideband frequencies. Note however, that if the desired signals were only 10 dB above the noise floor, the dirty LO sideband spurs would be well below the receiver noise floor, and therefore unimportant.

LO with Noise Sidebands

If the LO has phase noise, it may have a spectrum similar to figure A5. Now the clean signals are modified by the noise sidebands. If a third, weaker signal is between the two strong signals, it is completely obscured by the LO noise sidebands on the stronger signals. This is usually referred to as reciprocal mixing, and is the reason that good LO phase noise is so important in HF receivers for crowded bands.

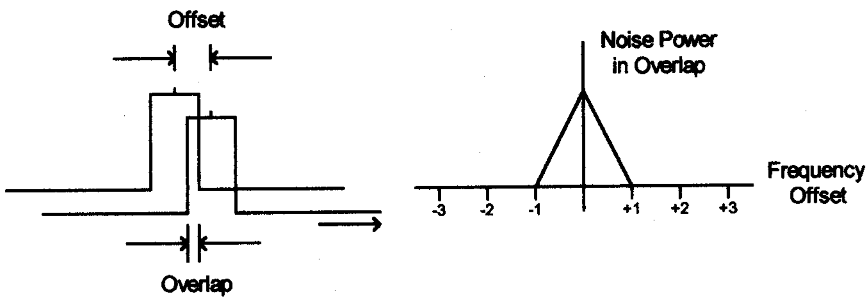

Figure A1

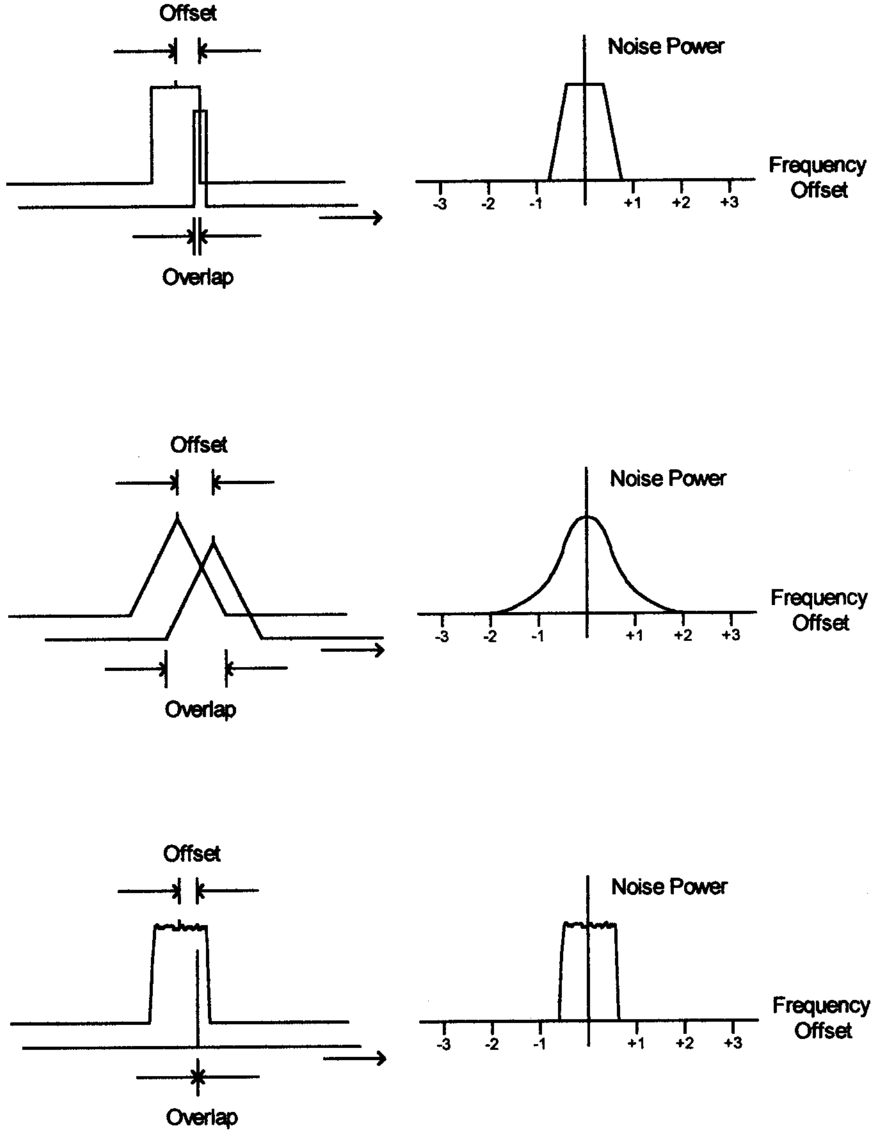
Offset
Overlap
Noise Power
Frequency Offset
-3
-2
-1
+1
+2
+3
Offset
Overlap
Noise Power
Frequency Offset
-3
-2
-1
+1
+2
+3
Offset
Overlap
Noise Power
Frequency Offset
-3
-2
-1
+1
+2
+3

Figure A2

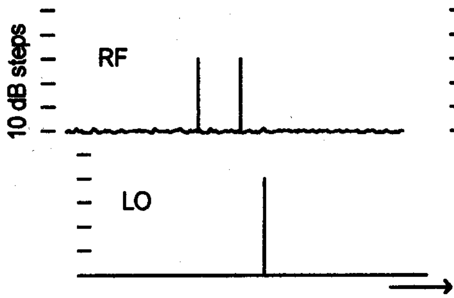

Figure A3

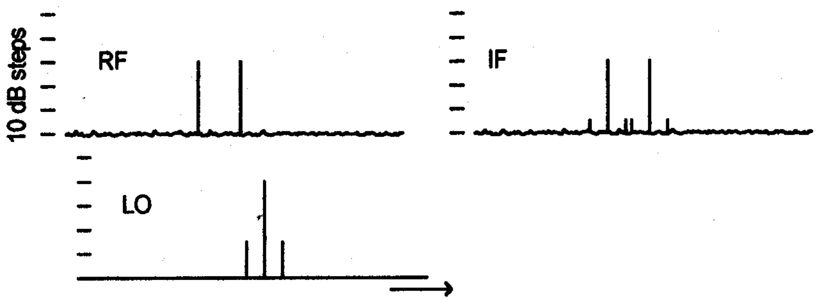

Figure A4

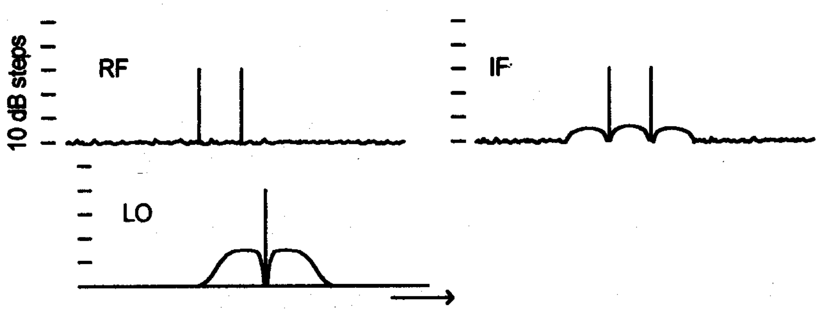

Figure A5