SARL HF Transceiver Project
Analog-to-Digital Converter Considerations

Introduction

The receiver architecture proposed for the SARL HF Transceiver project is a direct-conversion receiver with image rejection using digital signal processing (DSP)[1]. In order to be processed by the DSP, the signal must first be converted from analog to digital using an analog-to-digital converter (ADC). This paper examines the impact of the analog to digital conversion process on the system performance. It is not concerned with the limitations of real-life ADCs, but rather the theoretical limitations that the sampling process places on system performance.

The effect of analog to digital conversion on system performance arises from two inherent features of the process:

1. The analog signal, which is continuous with time, is broken up into a limited number of periodic values by the ADC, a process known as *sampling*;

2. While an infinite range of signal levels can be represented by an analog signal[2], the ADC represents the signal level at a point in time by one of a finite number of possible values, a process known as *quantization*.

In this paper I will first look at the effects of sampling and quantization separately, and then show how together they affect system performance.

The Effects of Sampling

The Nyquist sampling theorem states that any signal with a finite bandwidth B can be represented by, and uniquely reconstructed from, periodic samples of the signal provided that

$$B < S/2$$

where B is the bandwidth of the signal in Hz and S is the sample rate in samples per second. Conversely, if the bandwidth $B \geq S/2$, then the signal cannot be uniquely reconstructed from periodic samples taken at the sampling rate of S samples per second.

The reason why a signal with excessive bandwidth cannot be uniquely reconstructed from samples, is that there are actually an infinite number of signals that can be represented by any set of samples. If the bandwidth of the original signal is not limited as required by the Nyquist sampling theorem, then there is ambiguity as to which of the possible signals that could be represented by a particular set of samples is the actual signal that was sampled.

This is illustrated in Figure 1, which shows four sine waves of different frequencies. Assume that any one of these signals is sampled at the discrete time points represented by the numbers on the X axis, and the dashed grid lines. The samples will be identical, no matter which of the signals is sampled, since all the signals have identical values at the sample points. If the bandwidth of the original signal is unknown then it is impossible to reconstruct the signal from the samples, since one could not tell which of these possible signals was the correct one.

[1] Roos, Andrew (2007), *SARL HF Transceiver Project – Architecture Discussion Document.*
[2] In practice this is limited by noise; as we shall see, the effect of quantization can be modeled as additional *quantization noise*.

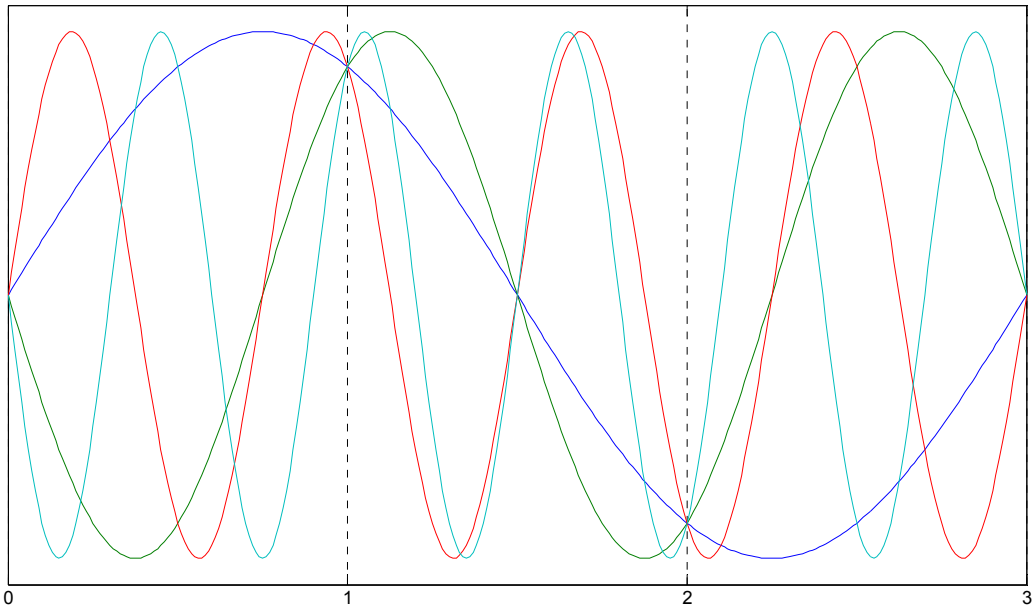

Figure 1 - Illustration of Aliasing

If the X axis is labeled in seconds, then the sampling rate is 1 sample per second. According to the Nyquist sampling theorem, we should be able to unambiguously reconstruct the original signal if its bandwidth is less than half the sampling rate, in this case less than 1/2 Hz. The frequencies of the signals shown in Figure 1 are 1/3 Hz, 2/3 Hz, 4/3 Hz and 5/3 Hz. Of these, only one frequency, 1/3 Hz, falls within the bandwidth limit of 1/2 Hz specified by the Nyquist sampling theorem, so if we know that the original sampled signal met the Nyquiest criterion, then there is no longer any ambiguity as to which is the correct signal.

In general, a signal of frequency F (where F < S/2) will have aliases at frequencies NS ± F, where S is the sampling rate and N any natural number (1, 2, 3, 4, ...).So in this example, where the sampling rate S = 1 sample per second, and the signal frequency F is 1/3 Hz, the alias frequencies are N.1 ± 1/3. For N=1 this gives aliases at frequencies 2/3 Hz and 4/3 Hz; for N=2, at 5/3 Hz and 7/3 Hz, and so on. The lowest frequency alias is found when N=1, at frequency

$$F_{alias} = S - F_{signal}$$

where S is the sampling rate and Fsignal the frequency of the sampled signal.

In order to eliminate problems caused by aliasing, practical systems typically include an “anti-aliasing” low-pass filter in the signal chain before the ADC. The purpose of this filter is to attenuate potential alias signals sufficiently to prevent them from affecting the processing and reconstruction of the desired signal. What constitutes “sufficient” attenuation depends on the application and is a question to which we shall return later. Note that if the bandwidth of the signal to be digitized is B Hz, and the sample rate S samples per second, then the lowest frequency signal that the anti-aliasing filter needs to attenuate is S-B Hz. For example, if the bandwidth is 2400 Hz and the sampling rate is 6000 samples per second, then the anti-aliasing low-pass filter must attenuate signals with frequencies from 6000-2400 = 3600 Hz upwards. This may require a fairly sharp filter. On the other hand if the same bandwidth signal is sampled at 48000 samples per second, then the lowest frequency alias signal is at 48000-2400 = 45600 Hz, which makes the design of the anti aliasing filter much simpler, in exchange for needing a faster ADC and DSP.

The Effects of Quantization

The signal level (for example, voltage or current) of an analog signal can take on an infinite number of possible values. For example, a voltage could be 0.1 V, or 0.11 V, or 0.111 V ... or 0.1111111111 V, and so on. However the output from an ADC is typically limited to a certain number of binary bits per sample. For example, an 8-bit ADC represents each sample by an 8-bit binary number, which means that there are only $2^8 = 256$ possible "levels" in the digitized signal. The process of assigning a discrete value to an analog signal is known as "quantization".

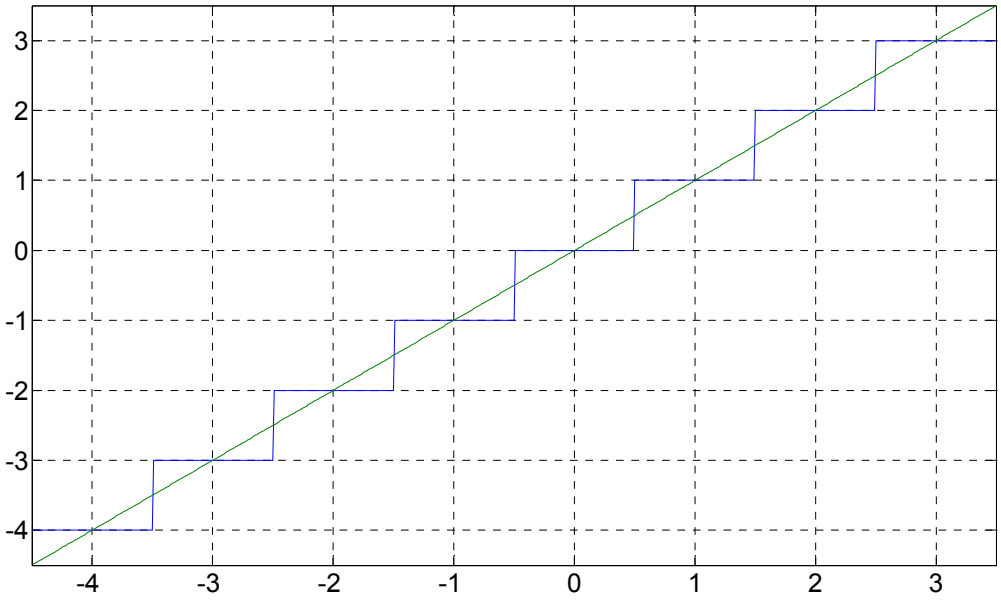

Figure 2 – Quantization

Figure 2 illustrates the effect of Quantization using a hypothetical 3-bit ADC with $2^3 = 8$ possible output values. The X axis represents the input voltage to the ADC, although the scale is unimportant (you can imagine it to be in Volts, if you would like). The Y axis represents the output of the ADC. Two lines are shown. The straight line shows what the output of the ADC would be if it did not quantize the input voltage – i.e. a "perfect" ADC with an infinite number of possible output values. Because it has an infinite number of output values, it can exactly represent any possible value of the input signal, and does not add any noise or distortion to the input signal. The stepped line shows the output of our hypothetical 3-bit ADC, with 8 possible output levels from -4 to +3. Since it cannot represent all possible input voltages exactly, it "rounds off" the input voltage to the nearest whole number, distorting it somewhat.

The output of the 3-bit ADC can be thought of as the output of the "perfect" ADC, plus an unwanted noise signal that represents the difference between the "perfect" ADC and the 3-bit ADC. The level of this "noise" is between -0.5 and +0.5 and it varies systematically with the input voltage as shown in Figure 3. This noise, which is an unavoidable consequence of the quantization process, is called "quantization noise". The maximum value of the quantization noise is one half of the voltage change represented by the least significant bit of the ADC. However when we calculate signal to noise ratios and such, it is common practice to talk in terms of root mean square (RMS) voltages, not peak voltages. Mathematically it can be shown that the RMS

voltage of the quantization noise is 1/√12 (about 0.29) of the voltage change represented by the least significant bit of the ADC[3]. In decibel terms, this is

$$20 \log10 (1/\sqrt{12}) = -10.8 \text{ dB}$$

i.e. the RMS noise power is 10.8 dB below the incremental power represented by the least significant bit of the ADC[4].

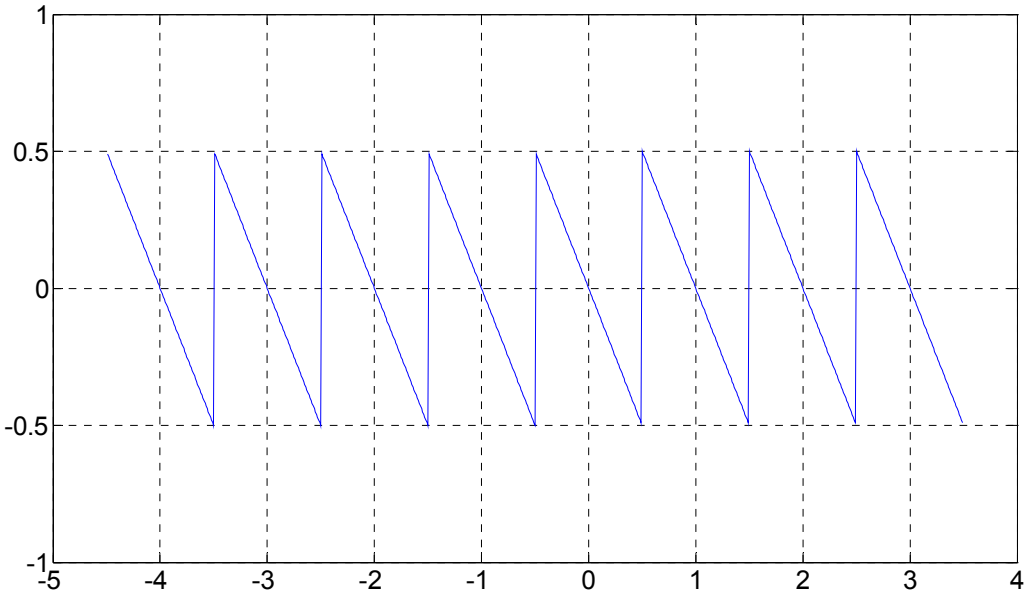

Figure 3 - Quantization Noise for a 3-bit ADC

This useful result allows us to calculate the theoretical signal-to-noise ratio (SNR) for an ADC. If the input signal is a sine wave, scaled so that the maximum and minimum peaks of the sine wave correspond to the maximum and minimum output values of the ADC, then the peak-to-peak level of the sine wave is 2^N least significant bits (LSBs) of the ADC, where N is the number of bits precision of the ADC. The RMS value of this sine wave is 1 / (2 √2) of the peak-to-peak value, or 9.0 dB below the peak value. The RMS noise level, as we have seen, is 10.8 dB below the power represented by the LSB. So the signal to noise ratio for a full-scale sine wave is

$$20 \log_{10} (2^N) - 9.0 + 10.8 = 6N + 1.8 \text{ dB}$$

where N is the number of bits in the ADC. For example, an 8-bit ADC (N=8) would have a theoretical signal-to-noise ratio of 6 * 8 + 1.8 = 49.8 dB.

Since an ADC will distort a signal badly by *clipping* it if the input voltage exceeds the input range of the ADC, it is important to ensure that the *peak* of the input signal is still within the range of allowable input voltages for the ADC. In typical (uncompressed) speech signals, the ratio of the peak-to-peak value of the signal to the RMS value is about 12 dB. Hence the signal-to-noise ratio available without clipping is

---

[3] Oppenheim, A and Schafer, R (1998) *Discrete-Time Signal Processing*, p. 196. Prentice Hall, New Jersey.

[4] Remember that decibels represent ratios of the signal *power*. Hence when dealing with voltages, the formula for decibels is 20 $\log_{10}$ (V1/V2) as power is a function of voltage *squared* provided the impedances are the same.

$$20 \log_{10} (2^N) - 12 + 10.8 = 6N - 1.2 \text{ dB}$$

So our 8-bit ADC would have a maximum SNR of 8 * 6 – 1.2 = 46.8 dB with speech signals.

Over-sampling

Over-sampling means sampling a signal at a rate faster than required by the Nyquist sampling theorem (i.e. faster than 2 B samples per second, where B is the bandwidth in Hz). This has two benefits. The first is that it moves the lowest alias frequency (S – B) away from the highest signal frequency (B), making it easier (and perhaps cheaper) to implement a suitable anti-aliasing filter. With sufficient over-sampling, the anti-aliasing filter might be no more than a single-section RC filter.

Although over-sampling can be used to reduce the requirement on the anti-aliasing filter, we must remember that in this receiver architecture, the anti-aliasing filter also serves as a roofing filter to prevent strong signals that are some distance from the desired frequency from overloading the ADC or causing inter-modulation distortion (IMD) in the base-band gain stages. Hence, despite the potential benefits of over-sampling, we still need to use reasonably sharp low-pass filters in the base-band signal path. Low-pass filters with a cut-off frequency of 2.7 kHz and a stop-band starting at about 5 kHz might be a good compromise between filter complexity and receiver performance. This suggests a minimum sampling rate of 7.7 kHz (2.7 kHz + 5 kHz) to avoid aliasing problems. In practice, a sampling rate of about 12 kHz would allow 2x over-sampling with an easily achievable shape factor of 1.1 for the decimation filter in DSP. In this case the analog low-pass filter stop-band must start not higher than (12 kHz – 2.7 kHz) 9.3 kHz to prevent aliasing.

The second benefit of over-sampling is that it turns out (although I won't attempt to show here) that the quantization noise is distributed evenly over the entire bandwidth of the ADC, up to the Nyquist limit S/2. Consider a signal with a bandwidth of 3 kHz being sampled by an ADC at the minimum sampling rate of 6 ksps. There is a certain amount of quantization noise introduced by the ADC as we have seen. This noise is spread over the 3 kHz bandwidth of the ADC. Now suppose we sample instead at 24 ksps. The *same* quantization noise power is generated, but it is spread evenly over a wider bandwidth, 12 kHz. We can now use digital filtering to reduce the bandwidth to 3 kHz, and the filter will filter out the quantization noise between 3 kHz and 12 kHz, which is ¾ of the quantization noise! So by sampling 4 times as fast as necessary (called "4x over-sampling") we have achieved a 6 dB reduction in the quantization noise. Another way of looking at it is that by 4x over-sampling we have reduced the quantization noise by the same amount that adding one extra bit to the ADC would have.

There are two limits to the use of over-sampling to increase the effective number of bits of an ADC. The first is that since only the quantization noise is reduced by over-sampling, when the noise contribution of other noise sources (such as the input voltage noise of the ADC) reach a similar magnitude to the quantization noise, then the total ADC noise is no longer dominated by quantization noise and additional over-sampling is ineffective. Secondly, the over-sampling required is an exponential function of the number of additional "effective bits" required for the ADC, so it becomes impractical to add to many "extra" bits worth of noise reduction through over-sampling. For example, to reduce the noise factor of an ADC by an amount equivalent to adding an extra 8 bits, one would have to over-sample by a factor of $4^8 = 65\,536$.

Dynamic Range

"Dynamic Range" (DR) means roughly the ratio between the strongest and the weakest signals that a system (in our case, the receiver) can successfully process. It is not a single performance indicator; in fact, there are many different "dynamic ranges" that are important in a receiver. The one that is most dependant upon the ADC characteristics is the two-signal, close-in dynamic

range. In this scenario, there is a weak signal that you are trying to listen to, and another much stringer signal nearby. The undesired strong signal is so close that it passes through all analogue filtering and arrives unattenuated at the ADC, along with the desired weak signal. Since both are present simultaneously at the ADC, we must rely on the dynamic range of the ADC in order to be able to digitize both signals successfully. Once digitized, the undesired signal can be removed by digital filtering; however we must ensure that its mere presence does not prevent the weak signal from being properly digitized by the ADC.

If we assume for the moment that the receiver has some sort of analog AGC (more on this later), then the problem is not so much that the strong signal may cause clipping, but rather that the strong signal, by activating the AGC and reducing the gain for both signals ahead of the ADC, may cause the level of the weak signal to be reduced to the point where it is lost in noise – especially quantization noise, which we shall assume dominates the ADC noise.

Assume that the gain ahead of the ADC is tailored so that the quantization noise of the ADC corresponds to an input signal level of -130 dBm. Then assume that the level of the strong signal is S9 (50 µV into 50Ω), -73 dBm. Then the dynamic range required in the ADC is 57 dB. Allowing a further 3 dB headroom for the AGC system (to allow AGC overshoot of 3 dB without clipping in the ADC), we require a final ADC DR of at least 60 dB. Using the formula for DR with a voice waveform, DR = 6N – 1.2 dB, we can calculate that the effective number of bits required is 10.2 bits. This could be achieved with a 10-bit ADC using 2x over-sampling.

Practical Considerations

Up to now, the paper has dealt with the theoretical performance limitations on ADCs. However like most components, real-world ADCs only approximate the theoretical. In particular, a combination of voltage noise at the input to the ADC, and non-linearity within the ADC, increases the noise floor of the ADC above that to be expected purely from Quantization noise. A good rule of thumb is that the additional noise and distortion found in a real-world ADC generally reduces its dynamic range by about 6 dB over the theoretical achievable. This is equivalent to having one less bit than it really has; we can say that the *effective number of bits* for a real-world ADC is generally about 1 bit less than the actual number of bits.

Hence to get the minimum performance required, which is that of a theoretical 10-bit ADC with 2x over-sampling, we would in fact need an 11-bit ADC with 2x over-sampling. We could also still use a 10-bit ADC with an extra 4x over-sampling to make up for the “lost” bit. When combined with the original 2x oversampling, this suggests a 10-bit ADC with 8x over-sampling as being about the minimum practical arrangement.

Looking around at real-world devices, we find that the dsPIC30Fxxxx series includes devices with both 10-bit and 12-bit ADCs. Although we have seen that a 10-bit ADC with 8x over-sampling would give us the minimum performance required, the additional dynamic range achievable with a 12-bit ADC would certainly be welcome as it would give us some “engineering margin” over the minimum specification. Using 12-bit ADCs would also allow us to use less over-sampling, reducing the required throughput of the ADC. I therefore suggest that we consider using a DSP with 12-bit ADC, and 4x over-sampling.

Specification sheets for ADCs may quote the “Signal to Noise And Distortion” (SINAD) for the ADC, and/or the “Effective Number of Bits” of the ADC. The SINAD is simply the practical dynamic range *for a full-scale sine wave input* of the ADC. Note that for voice signals, the actual achievable dynamic range is at least 3 dB less than the dynamic range quoted for a pure sine wave. The “Effective Number of Bits” is the number of bits that a perfect ADC would have in order to have the same dynamic range as the real-world device. Since noise and distortion always reduce the dynamic range of real-world devices, the effective number of bits is always less than the actual number of bits. Note that the effective number of bits may be non-integer.

For example the on-board ADCs on the dsPIC30F4013 specifies a SINAD of 68 dB (typical) and an effective number of bits of 10.95 (minimum) to 11.1 (typical). A quick check shows that with 11.1 bits, we would expect a dynamic range for a full-scale sine wave input of

$$6N + 1.8 = 68.4 \text{ dB}$$

which tallies with the quoted SINAD figure. These two figures are affectively telling us the same thing. Using the minimum effective number of bits (10.95), we can calculate the dynamic range for voice signals as

$$6N - 1.2 = 64.5 \text{ dB}$$

4x over-sampling will add 6 dB to the DR, so we gat a total of 70.5 dB DR for the practical 12-bit ADC with 4x over-sampling. Allowing 3 dB for AGC headroom, we end up with a predicted two-signal close-in DR of 67.5 dB, 10 dB above the minimum requirement.

One reason for choosing 4x over-sampling is to make the design of the low-pass filters in the base-band signal path practical. If B is the signal bandwidth (in Hz), then at 4x over-sampling, the sample rate will be S = 8B, and the lowest-frequency alias signal is at S-B, or 7B.

In practical terms, the base-band signal bandwidth will probably be in the region of 1.4 kHz, which will allow a 2.8 kHz audio bandwidth[5]. This means that the stop-band of the low-pass filter must start at 7 * 1.4 kHz = 9.8 kHz. This will give good protection from strong interfering signals 10 kHz or more away from the desired signal. This answers the question of what is considered a "close-in" signal in the term "close-in dynamic range" – the answer any signal within 10 kHz of the desired signal

Automatic Gain Control

The dynamic range considered so far is the two-signal, close-in dynamic range. It assumes that there are two signals entering the receiver at the same time. IN the analysis I assumed that we needed to cater for the case where one signal was S9 and the other was at the noise floor.

However there is another case where dynamic range is important. This is when the two signals occur separately, so we are only receiving a single signal at a time. This is what I call the "single-signal dynamic range" and it is the difference between the strongest signal that the receiver can receive, and the weakest signal that the receiver can receive.

Bearing in mind that the strongest signal might be a nearby transmitter, it is reasonable to expect to be able receive signals up to say S9 + 40 dB (-33 dBm) without the DAC clipping. The weakest signal requirement remains the same, only a bit above the ADC noise floor of -130 dBm. This requires a single-signal dynamic range of 97 dB.

To highlight the difference between the single-signal and two-signal dynamic range figures. I am saying that in the presence of a nearby S9 signal, an S1 signal should still be readable; and that an S9 + 40 dB signal should be readable on its own, although the presence of the S9 + 40 dB signal may mask a nearby S1 signal.

[5] Using the Weaver method of SSB demodulation, the two base-band channels (I and Q) together carry twice the bandwidth of any single channel, so the two 1.4 kHz baseband channels would give us a total audio bandwidth of 2.8 kHz. This method is recommended as it keeps the filter bandwidth to a minimum, making it easier to attenuate strong nearby signals, as well as reducing the required sample rate of the ADC and hence the DSP throughput requirements.

There are two ways to achieve the single-signal dynamic range requirement. The first is to increase the number of bits in the ADC to provide the required 97 dB dynamic range. This would require

$$(97 + 1.2) / 6 = 16.4 \text{ effective bits}$$

If we assume that a real-world 16-bit ADC has 15 effective bits, then we need an extra 1.4 bits which we can get by oversampling by a factor of

$$4^{1.4} = 7.0$$

In practice 8x over-sampling might be used, giving a dynamic range of 97.8 dB. However this solution would require great attention to be paid to the design of the baseband gain stages, to ensure that they maintain the low noise figure required to benefit from a 16-bit ADC, as well as the strong signal-handling capability to maintain the required dynamic range. Not to mention the gain and phase matching required to ensure adequate opposite-sideband rejection. There would also be additional costs associated with the off-board CODEC required to get 16-bit ADCs.

A better solution would be to provide 30-40 dB of gain control in the analog signal chain, and allow the DSP to adjust the gain to prevent the ADCs from clipping. This reduces the requirements on the base-band gain stages.

Conclusion

This analysis suggests that 12-bit ADCs with 4x over-sampling will give a two-signal close-in dynamic range of approximately 67 dB, which is 10 dB better than the minimum requirement. This can be combined with 30-40 dB of analog AGC action to give the required single-signal dynamic range of at least 97 dB.

If 4x over-sampling is used, then the base-band low-pass filters should have a pass-band cut-off frequency of 1.4 kHz, to give an audio bandwidth of 2.8 kHz, and a stop-band starting not later than 9.8 kHz, in order to avoid aliasing. The filter stop-band requirement is a minimum of 60 dB (preferably 80 dB) to meet the “spurious and image rejection” requirements in the specification. This requires a cut-off of 21.4 to 28.5 dB per octave, which should be achievable with a 4-5 pole filter. The filter requirement could be eased by using 8x over-sampling if necessary.

Andrew Roos ZS6AA
June 2007

Glossary

ADC	Analog to Digital Converter
DR	Dynamic range, the difference (usually expressed in decibels) between the strongest and weakest signals a system can process.
DSP	Digital Signal Processor/Processing
LSB	Least Significant Bit
quantization	Approximating the value of an analog signal by assigning it one of a finite number of possible values.
RMS	root-mean-square
sampling	Using periodic samples of the instantaneous value of a signal to represent the signal.
SINAD	Ratio of Signal to Noise and Distortion.