Sampling around Nyquist holes in high-speed converters

Luke Allen
Application Engineer
High-Speed Converters

Rob Reeder
Application Engineer
High-Speed Converters

Chase Wood
Application Engineer
High-Speed Converters

Frequency planning is a paramount part of any frequency-based application. Making sure that the frequency or frequencies of interest are valid, within band, and not reduced to the point where you lose the spurious dynamic range should be part of any frequency development strategy. In this paper, we will cover high-speed analog-to-digital converter (ADC) Nyquist rules as they apply to super Nyquist sampling. We will also cover how to use decimation to aid in your frequency planning and how to guard band against falling into a frequency "hole" during the design and development stage.

Nyquist rules

Super Nyquist sampling, intermediate frequency (IF) sampling and subsampling are popular in many frequency-based applications that employ software-defined radio (SDR) or radar-like receiver architectures (see **Figure 1**).

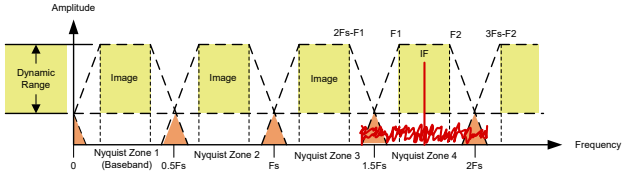

Figure 1. *Example of a super Nyquist sample vs. sampling at baseband (first Nyquist).*

There are two main reasons to plan frequencies outside of baseband (first Nyquist). The first reason is to gain the relaxation constraints put on the antialiasing filter design (AAF) (see **Figure 2**). Initially, the general filter rolloff needs to be much steeper when designing a baseband filter versus a filter design for a higher Nyquist zone. A steeper filter rolloff leads to a more complex filter where passive components become cumbersome. It's simple physics; you cannot purchase a 100µH inductor in a 0201 size. Therefore, when employing a higher Nyquist zone and possibly a higher sampling rate, the trade-offs and requirements for the rolloff in the stopband region are more relaxed, resulting in fewer components and smaller component sizes.

The second reason to use the high-frequency subsampling technique is to relax the radio-frequency (RF) receiver signal chain in front of the ADC. Assuming that the ADC can support the bandwidth requirements beyond the first Nyquist, which is almost always the case, relaxing the receiver signal chain could eliminate one or even two mix-down stages in the RF signal chain, resulting in even fewer components, less noise and less complexity.

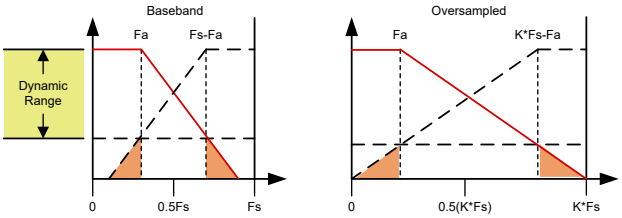

Figure 2. *Dynamic range vs. AAF stop-band attenuation.*

For example, **Figure 3** shows the Texas Instruments (TI) ADC3669 sampling an intermediate frequency of 800MHz relative to a 500 MSPS sampling frequency (Fs). Essentially, the signal is in the fourth Nyquist zone. The image or alias of the interest frequency reflects back to the first Nyquist zone appearing as a 200MHz signal. Most fast Fourier transform (FFT) analyzers, such as **High-Speed Data Converter Pro**, only plot an FFT of the first Nyquist zone, or 0Fs to 0.5Fs. Therefore, if the frequency of interest is above 0.5Fs, an image reflects down to the first Nyquist zone or baseband. This can make things confusing if spurious tones are in the band of interest as well.

So how does an ADC sample above 0.5Fs and still hold true to the Nyquist criteria? The Nyquist rule states that a signal must be sampled at a rate equal to or greater than twice its bandwidth in order to preserve all of the signal's information (see **Equation 1**):

$$\mathrm{Fs} > 2 > \mathrm{FBW} \qquad (1)$$

where Fs is the sample frequency and FBW is the maximum frequency of interest.

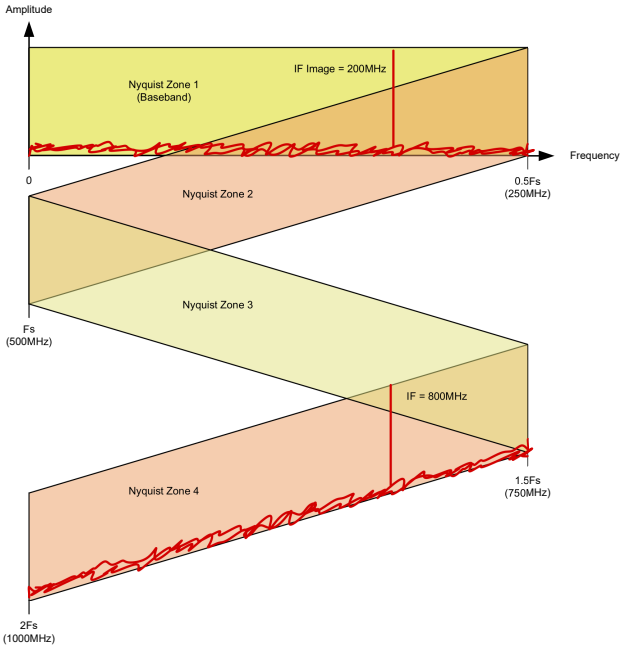

Figure 3. *ADC3669 example, where Fs = 500MSPS and the intermediate frequency = 800MHz.*

The key to holding the Nyquist rule to be true is the location of the frequencies of interest. As long as the signals do not overlap and stay within a single Nyquist zone, the Nyquist criteria still holds true. The only thing that has changed is the location of the first Nyquist zone to a higher one. IF sampling is becoming very popular because of these trade-offs.

What is process gain?

Earlier forms of high-speed signal chain line ups involved ADCs that used to consume watts of power per single channel, and FPGAs (field-programmable gate array) which captured, filtered and processed all converter data into some useful format. Most designers would use an approach called *process gain*. This approach not only helped in frequency planning by eliminating unwanted spurious and noise; it also enabled the ability to "gain" dynamic range in terms of the signal-to-noise ratio (SNR) by limiting the bandwidth processed within the Nyquist zone. Adding the process gain correction factor to the standard SNR equation results in **Equation 2**:

$$\mathrm{SNR} = 6.02 \times \mathrm{N} + 1.76\mathrm{dB} + 10 \times \log_{10}\left(\frac{\mathrm{Fs}}{(2 \times \mathrm{BW})}\right) \qquad (2)$$

where N is the number of ADC bits, Fs is the ADC sampling frequency and BW is the bandwidth of interest within the Nyquist zone.

With smaller process nodes deployed into both ADC and digital-to-analog converter technology, much of the standard FPGA digital features now reside within the ADC. Some examples include digital downconverters (DDCs), numerically controlled oscillators (NCOs) and frequency hopping. These features significantly help offload FPGA processing, making it possible to use its internal resources elsewhere.

Why frequency plan?

Proper frequency planning is an important step when designing systems using ADCs. Frequency planning ensures efficient utilization of the ADC's dynamic range and minimizes unwanted spurious signals, which are crucial for high-performance applications such as SDR systems or high-density RF signal chains.

One essential aspect of frequency planning is optimizing the ADC's dynamic range. Each additional signal in the analog domain consumes part of the ADC's available input signal power budget, reducing the total dynamic range capability. Proper frequency planning ensures that the ADC fully exercises its capabilities by strategically placing input frequencies to maximize the usable dynamic range. This approach involves considering where unwanted spurs and harmonics will appear in the sampled band, ensuring that there is no overlap between the signals of interest and spurious components.

Another critical aspect is managing the inherent spurious signals generated by ADCs, such as harmonics and interleaving artifacts. A frequency plan is successful when these spurious contributions remain out of the intended band, especially in systems that do not employ digital filtering after the data is sampled. Frequency planning also helps minimize the impact of clocking-related spurs, such as those that appear from clocking devices lacking channel isolation, which are particularly problematic in large element systems that use a high-density clock distribution solution.

Interleaving spur management is also a consideration in systems containing an interleaved ADC, where multiple ADC cores sample a common input signal orthogonally, effectively doubling the sampling rate and Nyquist region by a factor of two. This interleaving introduces spurious tones at Fs/2-Fin, however. Additionally, in systems that use higher interleaving factors, this same Fs/2-Fin spur is modulated once again by the "new" Fs/2-Fin. This results in the new Fin being comprised of the interleaved Fs/2-Fin spur, meaning that the spur count introduced is much higher than a single interleaving factor. Frequency planning enables you to mitigate these spurs by leveraging analog filtering or (preferably) digital decimation filters, attenuating them significantly without having to design a complex analog signal chain. While this approach reduces the instantaneous bandwidth to a selected region, it ensures better dynamic range and cleaner signal performance.

Common pitfalls in frequency planning

Despite its benefits, it is possible to have an improper frequency plan, which can lead to issues that degrade the ADC's performance. One common challenge is that the Nyquist zone overlaps. Poorly planned input signals may fall at the boundaries of Nyquist zones, leading to aliasing effects that reduce system performance. To prevent this, signals must be allocated within appropriate frequency bands to maintain spectral integrity with the Nyquist zone under consideration.

Clock spur contamination is another frequent issue, especially in systems using low-quality clocking devices or suboptimal clock distribution. These spurious signals, modulated onto the ADC spectrum, can severely impact sensitive applications by providing a known offset spur. Careful design of the clocking infrastructure, including the use of higher-quality clocking solutions, helps to mitigate these effects. Another possible method is to digitally filter the data with a high-order band rejection

filter at this offset frequency – although if implemented incorrectly, any wanted signal will be removed along with the spur.

Another challenge to overcome is the correction of tightly modulated third-order intermodulation distortion spurs. These spurs will almost always fall within the passband and are often the spur limiting the spurious-free dynamic range. In the event of a very high decimation factor, it might be possible that these tones land within the attenuation band. This is unlikely for most multitone systems, however, which inherently require a larger instantaneous bandwidth than single-tone systems and thus cannot incorporate such a large decimation filter.

Finally, you must navigate the trade-off between bandwidth and dynamic range. While decimation can suppress spurious components and harmonics, it comes at the cost of reduced instantaneous bandwidth. Balancing these trade-offs is essential to achieving the best performance for specific application requirements.

Advantages of proper frequency planning using decimation

Effective frequency planning delivers several benefits to enhance receiver system design. One advantage is improved spur suppression. Digital decimation filters effectively attenuate spurs, often achieving suppression levels around –85dBFS, which lead to cleaner signal performance and better utilization of the ADC's dynamic range for the intended signal rather than out-of-band spurious noise.

Another benefit is the reduction of data throughput from the ADC. By reducing the ADC's output data rate through decimation, you can interface the ADC with lower-speed, smaller and more cost-effective FPGAs. This reduction in transmitted data not only simplifies hardware requirements but also enables systems to operate in dual band or quad band, thus enabling the sampling of multiple RF bands simultaneously.

The ability of a system to be fully reconfigured in software alone is another significant advantage of using decimation on ADCs. You can plan the hardware interface between the ADC and FPGA to support the maximum data rate expected of the system, which would facilitate the ability to operate many other systems at lower data rates or more narrow bandwidths. Software-reconfigurable systems are particularly valuable in applications that require deployment into multiple scenarios.

Resource savings are also a notable outcome of effective frequency planning. By requiring fewer output lanes, whether high-speed serial data lanes or low-voltage differential signaling pairs, you can conserve valuable pins on both the ADC and FPGA, increasing the utilization factor. This is especially important in high-channel systems with printed circuit board area and power constraints.

Theoretical example: Frequency planning with decimation

Consider an interleaved ADC sampling a common RF input signal at Fs. The interleaving process introduces a spur at Fs/2-Fin, which can interfere with the intended signal. Applying a decimation-by-2 filter, as shown in **Figure 4**, makes it possible to attenuate this spur to levels within the decimation filter's rejection limits. Additionally, the decimation process reduces the ADC's output data rate, enabling cost-effective FPGA interfacing and simplifying downstream processing. Additionally, the wideband noise reduction introduces a 3dB process gain as a result of the N (that is Noise) in SNR being cut in half, while the S (that is Signal) stays the same.

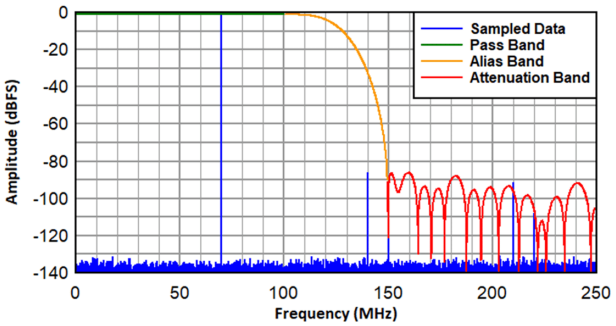

Figure 4. *Decimation filter response of theoretical data at 500MSPS with 70MHz FIN (with decimation factor of 2).*

Real World Examples: Frequency planning with decimation

The use of decimation on an ADC simplifies the frequency planning process because reducing the sampling rate effectively narrows the bandwidth of interest. Think of decimation as selectively focusing in on a narrower part of your spectrum. By focusing on a narrower band, more of the unwanted harmonics or spurs end up falling outside of the passband of interest, and in turn are filtered out. The following examples use the **ADC3669** to demonstrate the difference that decimation makes when frequency planning. **Figure 5** shows a traditional spectrum capture when the ADC is not performing decimation using an FFT size of 16384 points. You can see that the unwanted harmonics are in band and are negatively affecting performance.

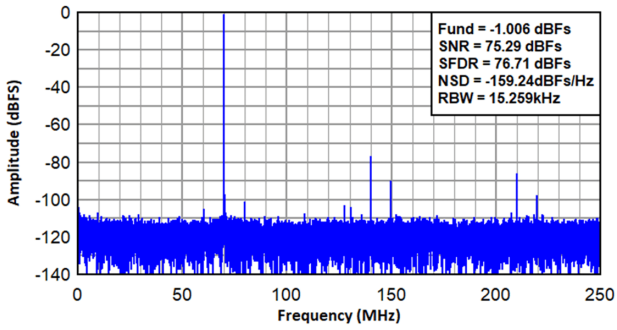

Figure 5. *Real Spectrum captured by the ADC3669 at 500MSPS with near 70MHz FIN (No Decimation).*

It could be that these harmonics are additive noise contributed by the ADC, or some external analog frequencies. **Figure 6** shows an example when the ADC is in real decimation mode, with a decimation factor of 2. You can see that the unwanted harmonic spurs now fall out of band and get filtered out by the decimation filter. Note that there is an additional +3dB improvement due to process gain.

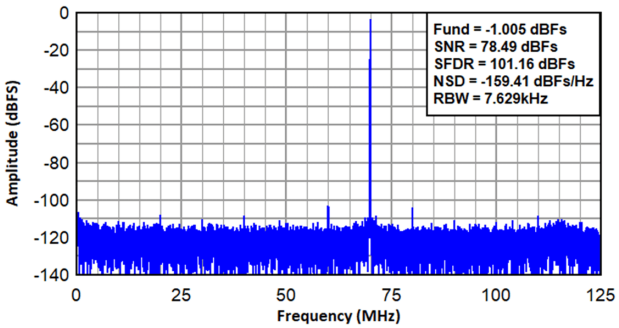

Figure 6. *Real Spectrum captured by the ADC3669 at 500MSPS with near 70MHz FIN (with decimation factor of 2).*

Additionally, the resolution bandwidth of the FFT actually reduces by a factor of two as well, since we maintain the same number of points for the FFT computation. This helps to resolve analog frequencies into closer bins. Up to this point, we have only talked about real decimation, which simply filters the data without any frequency shift. Real decimation is great if your signal of interest falls somewhere less than Fs/4 each time you decimate. But what if you want to decimate a signal that falls outside of this range? The signal of interest is often not centered at zero frequency (baseband), but rather at some intermediate frequency. This is where complex decimation comes in to play. ADCs with newer digital features, such as the ADC3669, incorporate an NCO mixer in the complex DDC stage. Mixing the signal of interest with an NCO frequency shifts the signal to baseband before decimation, enabling you to take advantage of the benefits of decimation for a signal anywhere within the device's bandwidth.

Figure 7 shows a capture of the ADC3669 in complex decimation mode with a decimation factor of 64 resulting in an effective sampling bandwidth of 7.8125MHz when the FFT is calculated using 8192 points. The input frequency is 70MHz and the NCO frequency is 71MHz. When the signal gets mixed with the NCO frequency,

the signal shifts to baseband, resulting in a tone at approximately -1MHz.

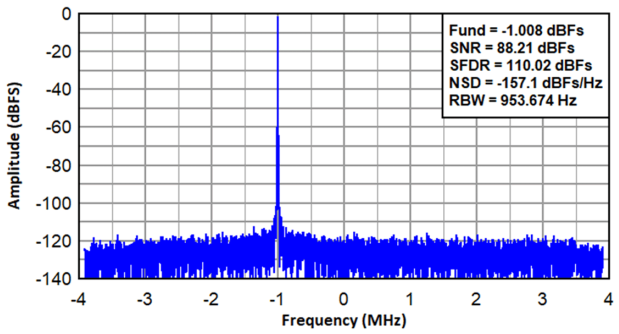

Figure 7. *Complex Spectrum captured by the ADC3669 at 500MSPS with decimation factor of 64 (FIN=70MHz, Fnco=71MHz).*

The ADC3669 can capture a narrow band with a decimation factor of up to 32768, which is beneficial for applications with high-density RF bands or tight channel spacing. Decimating by such a high factor allows you to zoom into your signal of interest, filtering out virtually everything else. The range of decimation factors offered by modern ADCs such as the ADC3669 enables more flexibility when frequency planning, as it is much easier to filter out unwanted spurs. **Figure 8** shows a capture with a decimation factor of 16384 calculated using 8192 FFT points, resulting in a resolution bandwidth of 3.726Hz. Even if your spurs are within thousands of hertz of the fundamental, you can easily filter them out with a high decimation rate.

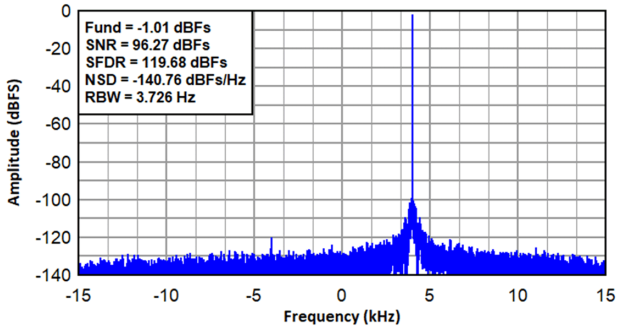

Figure 8. *Complex Spectrum captured by the ADC3669 at 500MSPS with decimation factor of 16384 (FIN=70MHz, Fnco=69.996MHz).*

As the NCO frequency is 4kHz lower than the input signal, the down converted signal appears at a positive frequency offset. While operating in this decimation mode and at 500MSPS, this ADC can sample signals within a 30.517kHz range around the programmable NCO frequency.

Conclusion

Frequency planning is an essential aspect of ADC-based system design, tackling challenges such as spur management, dynamic range optimization and AAF design, as well as efficient data handling. By implementing thoughtful frequency planning upfront, you can avoid common pitfalls such as Nyquist zone overlaps and clock spur contamination, while benefiting from advantages such as improved spurious suppression and dynamic range, a reduction in the ADC's digital interface or data rate, and FPGA resource savings. Carefully balancing these trade-offs and leveraging features such as ADC decimation make it possible to achieve high-performance, software reconfigurable receiver systems for a range of applications, while avoiding your next sampling hole.

Related Websites

- Discover techniques for fast frequency hopping in RF sampling data converters by checking out **SLYT861: Exploring fast frequency hopping in RF sampling data converters** by Chase Wood, featured in the Analog Design Journal, Oct 2024.
- Understand the relationship between sampling and data rate with **TIPL4701: Sampling vs. data rate, decimation (DDC) and interpolation (DUC) in high-speed data converters** by Jim Seton, published in Aug 2017.
- Gain insights into higher-bandwidth decimation examples by reading **Analyzing High-Bandwidth Spectrum Clusters** by Chase Wood, published in Embedded Computing Design, May 2024.
- Optimize your RF-sampling frequency planning using the **RF-Sampling Frequency Planner, Analog Filter, and DDC Calculator**. For detailed specifications, explore the **TI ADC3669 data sheet**.

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL AND RELIABILITY DATA (INCLUDING DATA SHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES “AS IS” AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS AND IMPLIED, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT OF THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for skilled developers designing with TI products. You are solely responsible for (1) selecting the appropriate TI products for your application, (2) designing, validating and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, regulatory or other requirements.

These resources are subject to change without notice. TI grants you permission to use these resources only for development of an application that uses the TI products described in the resource. Other reproduction and display of these resources is prohibited. No license is granted to any other TI intellectual property right or to any third party intellectual property right. TI disclaims responsibility for, and you will fully indemnify TI and its representatives against, any claims, damages, costs, losses, and liabilities arising out of your use of these resources.

TI’s products are provided subject to TI’s Terms of Sale or other applicable terms available either on ti.com or provided in conjunction with such TI products. TI’s provision of these resources does not expand or otherwise alter TI’s applicable warranties or warranty disclaimers for TI products.

TI objects to and rejects any additional or different terms you may have proposed.

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265