Application Note: Understanding the DSP Features in Plural™ Family of ADCs

Overview

Silanna's ADC portfolio includes a selection of feature rich parts that integrate DSP blocks with the 12-, 14- or 16-bit dual channel ADC. The DSP features include decimation, digital down conversion (DDC), and IQ mismatch correction. Many of our customers send the ADC output data to an FPGA for further processing. By including some of the commonly used DSP blocks as hard-wired logic in our ADC chip, we often enable the customer to use a smaller FPGA, reduce the FPGA power consumption, and reduce the data rate across the digital interface. Figure 1 shows the DSP signal path included in Silanna's feature rich ADC parts.

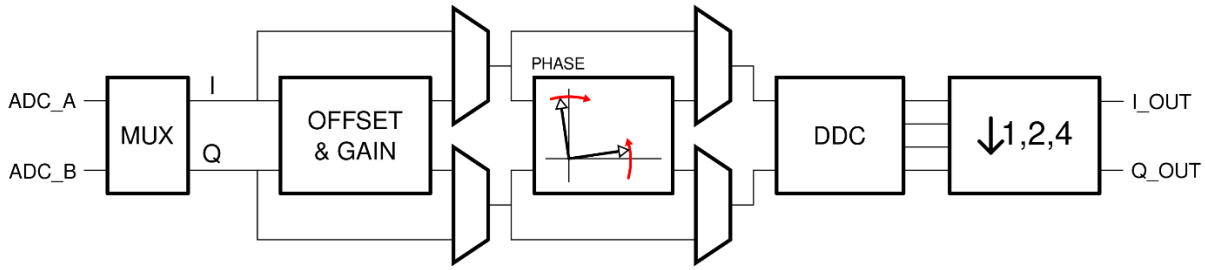

Figure 1. DSP path in Silanna's ADCs.

This application note describes the functions that can be implemented using the DSP blocks with more background information than our product datasheets do while leaving the details of the required register programming in the datasheets. For the benefit of those not completely familiar with the subject, we start the discussion with introduction to sampling effects, anti-alias filtering, and sub-sampling from higher Nyquist zones.

Effects of Sampling

An ADC is a device that converts an analog continuous time signal into a sampled digital one. This process can be divided into two, usually distinct, tasks: sampling and quantization. In a pipelined ADC the input voltage is sampled into a capacitor using a switch controlled by the sampling clock. This sampled and held voltage is then processed by the rest of the ADC into a digital word representing the magnitude of the sample.

The trajectory a continuous time signal takes between two adjacent samples is not captured by the sampling process, which can lead to loss of information. The Nyquist criterion states that a signal sampled using a sampling frequency (Fs) that is at least twice of its bandwidth, can be perfectly reconstructed afterwards. While this is a necessary condition for successful sampling it is not sufficient on its own. One must also consider the location of the frequency band of interest relative to the multiples of half sampling rate (Fs/2), which will be discussed later in this document.

The sampled domain (also known as discrete time domain) can represent signal frequencies from −Fs/2 to +Fs/2. When dealing with real valued signals (in the time domain) the negative frequency components are a mirror image of the positive ones and usually omitted from graphs and illustrations. The sampling process maps the infinitely long continuous frequency axis into this finite range leading into a phenomenon called aliasing where an infinite number of frequencies from the continuous time domain map into a single frequency in discrete time domain.

One way to visualize this mapping is to imagine the discrete time frequency axis as a circle (or a cross section of a cylinder) with the zero frequency at 3 o'clock position, +Fs/4 at 12, -Fs/4 at 6, and both +Fs/2 and −Fs/2 at 9 o'clock. The continuous time frequency axis is an infinitely long string or a thread with its zero frequency permanently attached to its counterpart on the discrete time circle. The frequency mapping is accomplished by winding the continuous time string around the discrete time cylinder. Following this process into positive direction up to Fs/2 shows one to one mapping from the continuous time domain into the discrete time domain. Beyond this point the frequencies on the string start to fall into the negative frequencies in the discrete time domain until at Fs the circle is completed. The next stretch of frequencies from Fs to 2Fs repeats the same sequence and so on, every loop mapping on top of the previous ones.

As mentioned earlier, when dealing with real valued signals (in time domain) the negative frequencies get reflected into the positive ones and vice versa. This means that a signal at continuous time frequency $Fs/2 + \Delta f$ folds (or aliases) on top of a signal at $Fs/2 - \Delta f$ because of the sampling. Therefore, while the bandwidth of signal of interest might be relatively narrow fulfilling the Nyquist criterion, it cannot span across a boundary defined an integer multiple of Fs/2. The frequency ranges between these boundaries are called Nyquist zones: the first being defined as the frequency range from 0 to Fs/2, the second from Fs/2 to Fs, the third from Fs to 3Fs/2, and so on. This gives the second condition for successful sampling: the signal must be entirely contained within a single Nyquist zone in the frequency domain.

The aliasing effect described above can be either a friend or a foe. It can be utilized for sub-sampling, but it also necessitates anti-alias filtering before the sampler.

Higher Nyquist Zone Sampling

Sub-sampling (or higher Nyquist zone sampling) is a setup where the analog input signal is in some Nyquist zone above Fs/2 and brought down to the first Nyquist zone by the sampling operation somewhat analogous to performing a down conversion with an analog mixer. This can benefit the system by eliminating a mixer or allowing the use of a lower speed ADC. One should notice that the spectrum of a sub-sampled signal appears flipped along the frequency axis when the input signal is in even Nyquist zones. Sampling a high frequency input signal requires good quality sampling clock because sensitivity to clock jitter increases with signal frequency.

Anti-Alias Filtering

In most real-world applications, the input signal to the ADC contains unwanted signal content outside the frequency band of interest. For instance, it can contain adjacent channels and neighboring transmission bands in the case of a wireless system. If the aliasing folds these frequencies on top of the desired signal band, they need to be suppressed to a level acceptable for the system performance using an anti-alias filter prior to sampling.

For instance, when operating in the first Nyquist zone, say we have a signal that extends from DC to 0.4*Fs. The anti-alias filter for this system is a low pass filter with pass band up to 0.4*Fs and stop band starting at 0.6*Fs, stop band attenuation being defined by expected worst case signal levels and system requirements. For instance, to achieve a −80dBFS SFDR while expecting −10dBFS unwanted tone level at 0.6Fs, the stop band attenuation specification must be set to 70dB. The complexity of the anti-alias filter is heavily influenced by the level of stop band attenuation and the width of the transition band. Besides the degree of the filter, the precision of the pole and zero locations get affected by both factors.

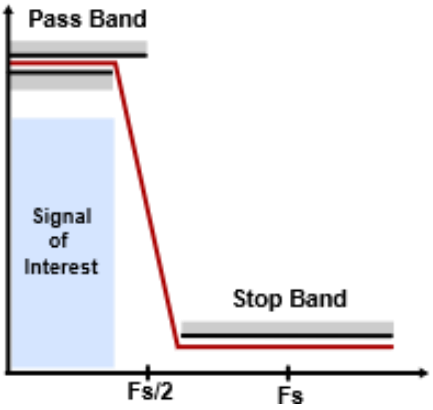

Figure 2. Low-pass anti alias filtering in 1st Nyquist zone.

When the input signal is in higher Nyquist zones the anti-alias filter must have a band-pass response. This, together with the fact that the complexity of the filter is proportional to the relative, not the absolute, width of the transition band, makes anti-aliasing increasingly difficult as the input frequency and the zone get higher. Nyquist gaps are the areas around multiples of Fs/2 that cannot be used for signal content because of impractically steep anti-alias filtering. In general, one should choose the sampling frequency of the system in a way that locates the signal band roughly in the center of a Nyquist zone and provides enough over-sampling for reasonable anti-alias filter transition bands. Going beyond the 3rd Nyquist zone is rarely practical.

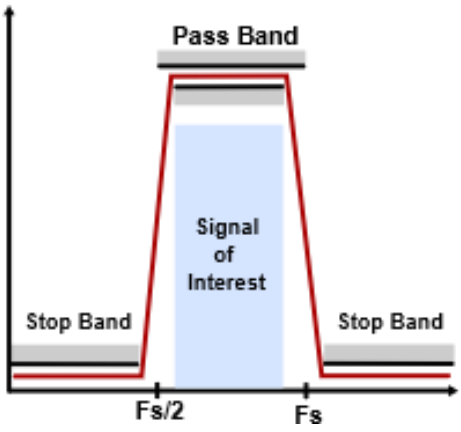

Figure 3. Band-pass anti alias filtering in 2nd Nyquist zone.

Over Sampling

Over sampling refers to the factor by which the sampling rate is made higher than required by Nyquist criterion.

ADCs show non-ideal effects that appear as white (flat) noise in the spectrum of the digital output signal. These effects are the thermal noise originating from the sampler and the other analog blocks in the ADC, quantization noise due to the finite resolution and accuracy of the converter, and wide band clock jitter. The output being a sampled signal means that all this noise is concentrated between DC and Fs/2. Using a higher sampling rate, i.e. increasing the oversampling factor, spreads this noise over more frequencies, making the noise density lower. Digital filtering can be used to improve the signal-to-noise ratio (SNR) by filtering out the noise outside the signal bandwidth.

Decimation

Decimation is a DSP technique that reduces the sampling rate of a digital signal. From the system efficiency point of view, it is desirable to represent the signal at the lowest rate suitable for the situation: transmitting a signal from point to point should have zero to minimal oversampling factor while an operation such as digital filtering benefits from having a modest over sampling factor. After analog to digital conversion the signal is commonly oversampled by a factor 1.5 or higher because of tradeoffs made for practical anti-alias filter design.

By utilizing the decimation feature included in our ADCs the user can offload some of the DSP burden from the FPGA to improve system cost and power efficiency. By selecting an ADC product with a built-in decimator, the user can keep the digital interface data rate low while sampling at a higher rate. This relaxes anti-alias filter specifications and improves the signal to noise ratio thanks to over sampling. The expected SNR improvement is about 3dB for every factor of two decimation. Figures 4-6 show measured ADC FFTs that demonstrate this effect.

The resampling operation suffers from the same aliasing effects as already discussed earlier in this document, hence the bandwidth of the desired signal must meet the Nyquist criterion of the new sampling rate. For that reason, the decimation is a two-step process: first the signal is digitally filtered to get rid of the spectral content at the frequencies that will alias on top of the defined signal band and then resampled at the new rate.

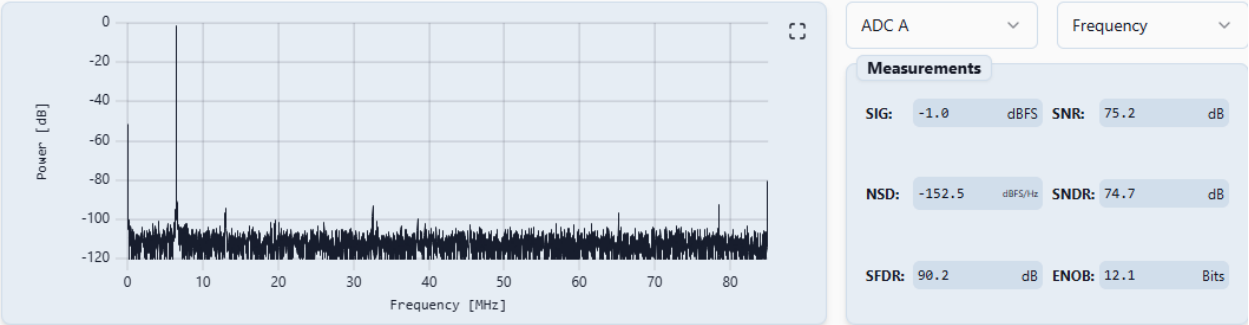

Figure 4. ADC sampling at 170MS/s.

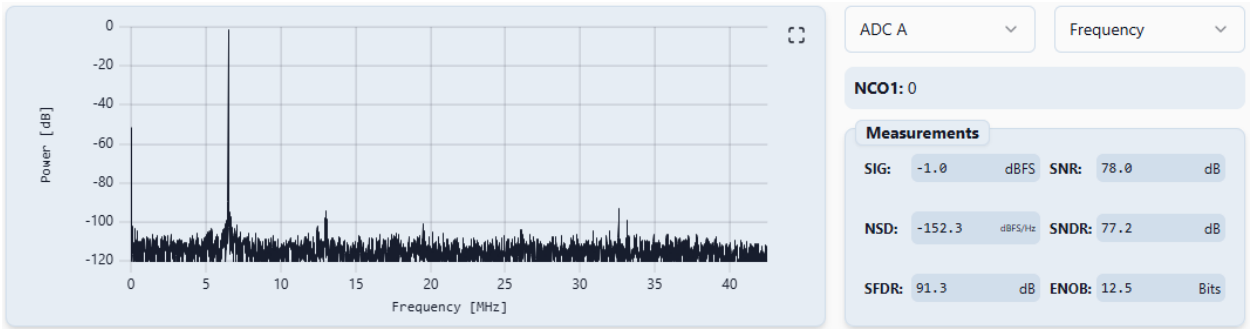

Figure 5. ADC sampling at 170MS/s, decimation by 2 yielding 2.8dB SNR gain.

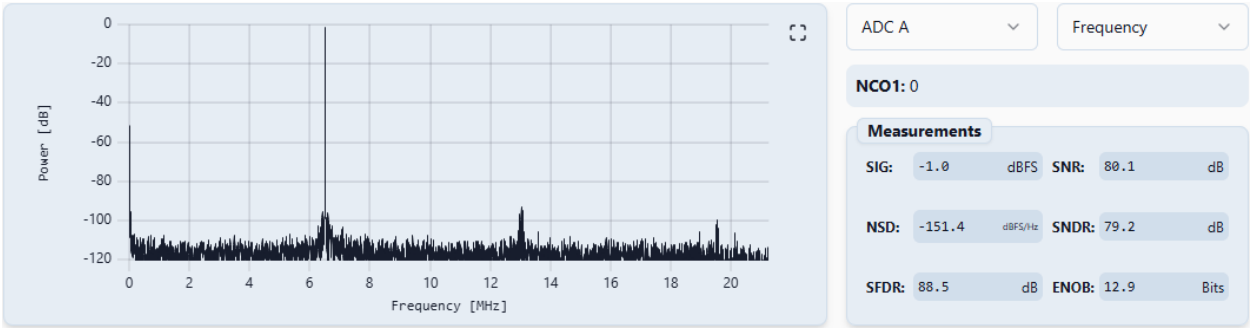

Figure 6. ADC sampling at 170MS/s, decimation by 4 yielding 4.9dB SNR gain.

For efficient hardware implementation it is customary to implement a high decimation ratio as cascade of several low ratio stages. For instance, a decimate-by-eight structure can be implemented as a cascade of three decimate-by-two stages. As in the case of analog filters, the complexity of a digital filter is proportional to the stop band attenuation and the steepness of the transition band. A typical value for the width of the filter pass band is around 0.4 times the final sampling rate, which makes the system 1.25 x oversampled.

Silanna's feature rich ADCs support decimating by factors 2 or 4 with 40% filter pass band and greater than 80dB stop band attenuation. The implementation uses FIR half-band filters with frequency responses shown in Figure 7. The filter taps are:

[-26, 0, 181, 0,-702, 0, 2056, 0,-5423, 0, 20299, 32767, 20299, 0,-5423, 0, 2056, 0,-702, 0, 181, 0,-26]

and

[-4, 0, 13, 0,-32, 0, 70, 0,-135, 0, 243, 0,-410, 0, 658, 0,-1015, 0, 1521, 0,-2236, 0, 3265, 0,-4829, 0, 7499, 0,-13384, 0, 41545, 65536, 41545, 0,-13384, 0, 7499, 0,-4829, 0, 3265, 0,-2236, 0, 1521, 0,-1015, 0, 658, 0,-410, 0, 243, 0,-135, 0, 70, 0,-32, 0, 13, 0,-4]

Decimation rate selection for two channels in dual ADC is common. Digital output data and clock are at the decimated rate while the sampling clock needs to be supplied at the ADC sampling rate (2x or 4x the output rate).

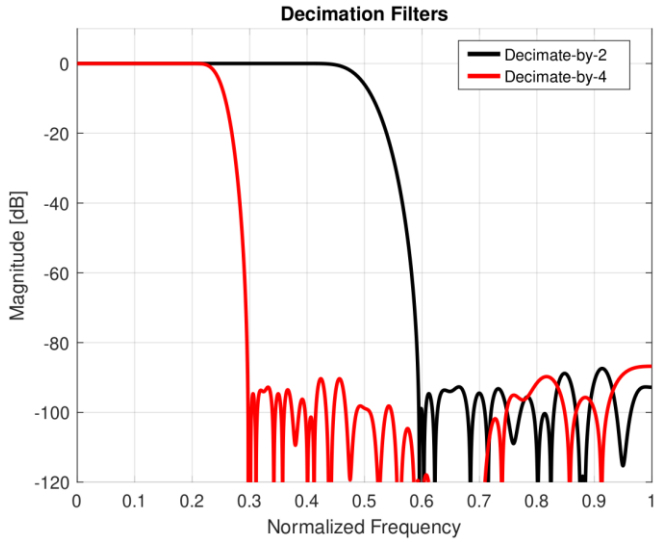

Figure 7. Decimator's input-referred frequency response.

Decimation, as described above, requires the signal band to be confined between DC and 0.2*Fs for decimation by two, or between DC and 0.1*Fs for a factor of four. Decimation for signals located in the upper half of the first or the lower half of the second Nyquist zone is enabled by adding a mixer in the signal chain for moving the signal to the pass band of the decimation filter. This mixer is a trivial circuit that multiplies the input samples with +1, -1, … or +1, +1, -1, -1, … sequence prior to the decimation. Using the first sequence is equivalent to a mixer with local oscillator (LO) at Fs/2 and the second one with LO at Fs/4. Using either sequence for signals below the LO (lower side band) flips the spectrum along the frequency axis. The use of the second sequence results in a 3dB loss of digital signal amplitude and adds 90° phase shift between the two ADC data because of implementation details.

Depending on the location of the signal band neither mixing down with Fs/4 or Fs/2 is optimal. The full featured DDC, described later in this document, can be used in those situations, with the difference that its digital output data is provided in IQ format.

Interleaving and Decimation

The two ADC cores in Silanna's feature rich dual ADC parts can be ganged together to operate as a single higher performance ADC. By wiring the ADC analog inputs together and providing 2x sampling clock the two ADCs can be time interleaved to operate as a single ADC with twice the sampling rate of a single core. When combined with 2x decimation the sample rate is reduced back to single core rate while the SNR typically shows 2.5 to 3dB improvement compared to a single core, which is illustrated by the ADC measurements shown in Figures 8-10. The interleaving feature is not supported without enabling the decimation.

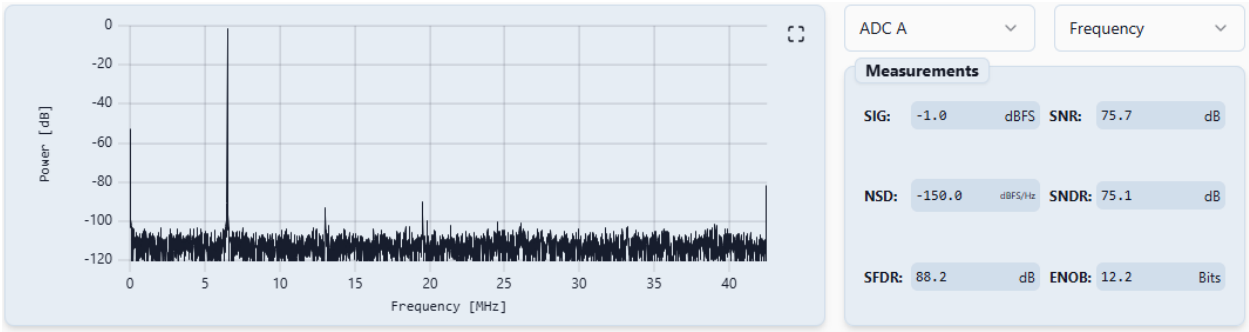

Figure 8. Single ADC core sampling at 85MS/s.

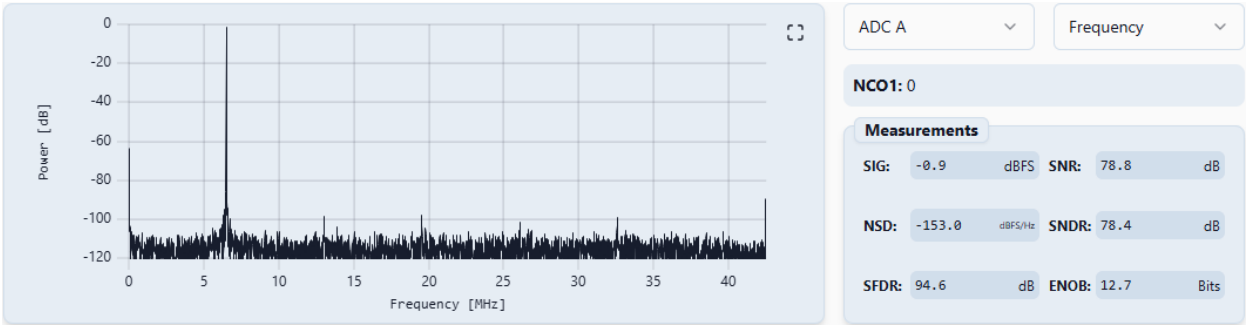

Figure 9. Two ADC cores sampling at 85MS/s, interleaved and decimated by 2, yielding 3.1dB SNR gain.

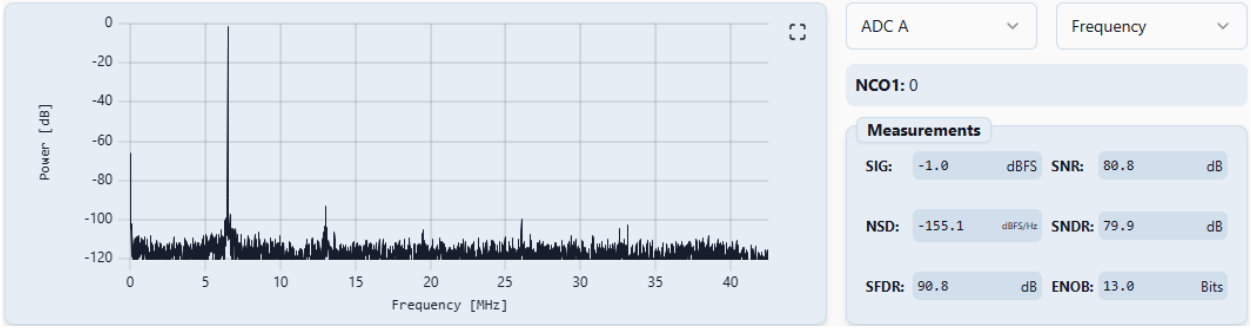

Figure 10. Two ADC cores sampling at 170MS/s, interleaved and decimated by 4, yielding 5.1dB SNR gain.

Besides improving the SNR, the interleave-decimate feature also relaxes anti-alias filter requirements. It should be noted, though, that interleaving artifacts due to mismatches between the two ADC cores and the external signal routing to the two ADCs are not calibrated out. Thanks to decimation, these artifacts do not cause self-interference but limit the amount the anti-alias filter rejection can be relaxed. The typical level of the interleaving artifacts is −40dBFS or better. This means that, for instance, if the system needs 85dB stop band attenuation from the anti-alias filter, the stop band needs to be −45dB from Fs/2 to 3Fs/2 and −85dB beyond that, as shown in Figure 11. A conventional ADC sampling at Fs needs attenuation of 85dB starting at Fs/2, which is substantially more demanding than the first case, which in a case study was handled with a 5th order filter while the latter demanded 9th order one.

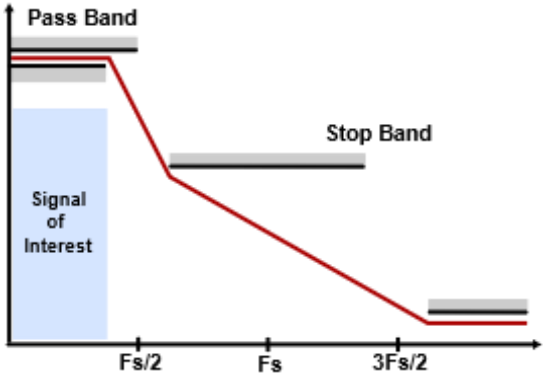

Figure 11. Interleaving and decimation relaxes anti alias filter specifications.

As discussed in the context of decimation (without interleaving) higher Nyquist zone sampling can be enabled by mixing down with a simple +1 / -1 sequence. The +1, -1, … sequence needed for LO at Fs/2 is simply a static signal inversion of one of the two ADCs. This can be done either in the analog domain by crossing the inputs one ADC or digitally in the DSP block. LO at Fs/4 is enabled by sequence +1, +1, -1, -1, …, which reduces digital output amplitude by 3dB.

In this mode the digital output is duplicated to the pins of both ADCs. The IOs of one of the ADCs can be programmed off if desired. The digital output clock is at half or one quarter of the sampling clock rate depending on selected decimation factor.

Good routing symmetry should be used when connecting the ADC inputs together.

Averaging

An alternative way to utilize a dual ADC as a single higher performance ADC, also supported by Silanna's feature rich parts, is to operate the two ADCs in parallel and average the digital outputs. This reduces the SNR noise components originating from thermal noise and quantization errors but doesn't significantly improve the noise originating from clock jitter because the two samples are captured on the same clock edge. When the system is not jitter-limited, SNR improvement from 2.5 to 3dB can be expected as demonstrated by measurements shown in Figures 12-13.

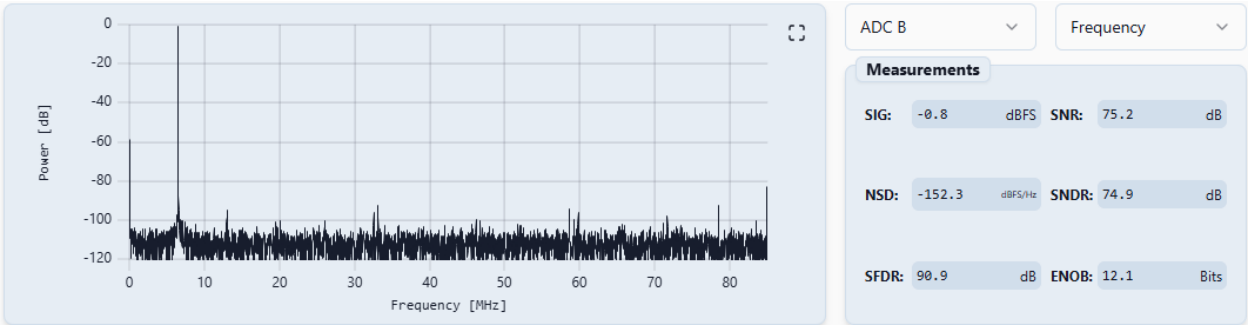

Figure 12. Single ADC sampled at 170MS/s.

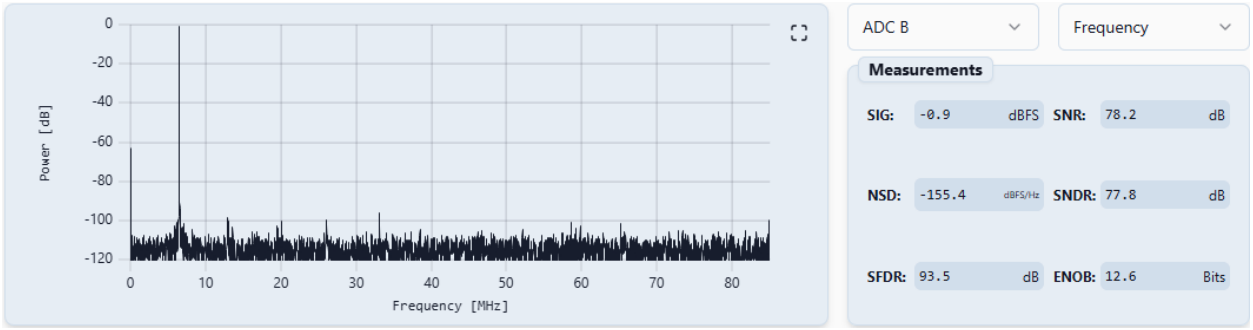

Figure 13. Two ADCs sampled at 170MS/s and averaged, yielding 3.0dB SNR gain.

In this mode the digital output is duplicated to the pins of both ADCs. The IOs of one of the ADCs can be programmed off, if desired. Driving the analog inputs of the averaged ADC pair can be more difficult than the interleaved one because both sampling capacitors are connected to the driver at the same time.

Digital Down Conversion

Representing signals in IQ format (complex values) has many benefits in DSP, such as being able to move signals around in frequency domain with simple mixing operations that don't require image rejection filtering and transforming band-pass filtering operations into low-pass ones. Centering the signal of interest at zero frequency often allows a reduction of the sampling rate which improves the system efficiency.

A bandpass signal sampled with an ADC can be converted into IQ representation and moved to zero center frequency using a digital down converter (DDC). It uses a numerically controlled oscillator (NCO) to produce a digital LO signal with sine and cosine components and two multipliers for implementing a digital mixer. The mixer output is an IQ signal at the input sampling rate.

In frequency domain this complex mixing operation rotates the signal spectrum around a circular frequency axis, which has its zero at 3 o'clock position and both +Fs/2 and −Fs/2 at 9 o'clock position. As a result, the bandpass signal of interest is moved clockwise to the vicinity of zero frequency and its mirror image in the negative frequencies rotated to higher negative frequencies, often crossing −Fs/2 and showing up partially at high positive frequencies as well.

As by creating I and Q signals from a real signal, we effectively doubled the sampling rate (by decoupling the negative frequencies from the positive ones), the IQ signal can be decimated by two without losing any unique spectral content. A decimation filter designed to match the bandwidth of the signal of interest removes the upconverted negative-frequency image. Applying a low-pass filter to I and Q components provides a symmetric band-pass response centered at zero frequency.

Silanna's feature-rich dual ADC parts implement independent DDCs for each ADC as shown in Figure 14. Decimation by two or four is supported after the digital mixer. The two DDCs can be programmed to use either a common NCO or independent NCOs. The decimation rate selection is common for the two units.

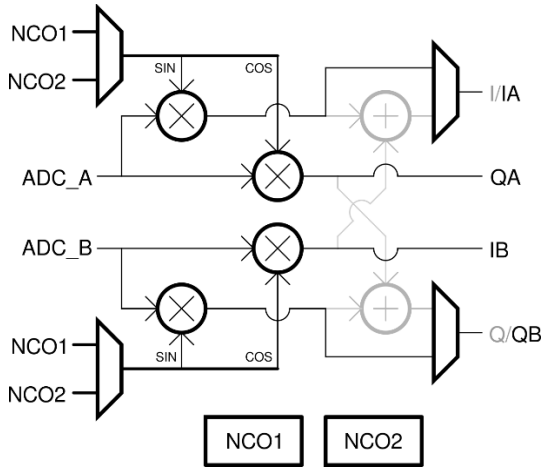

Figure 14. DDC configured for two independent signal paths.

The digital output IQ data is multiplexed to the ADC digital output pins, alternating samples representing I and Q respectively. The output pin normally used as overrange indicator can be programmed to serve as IQ indicator. A simple initial data synchronization step is needed on the receiver side. This is shown in detail in the product data sheet.

Using Silanna's ADCs with the DDC allows the user to reduce the DSP content on the FPGA. DDC and decimation reduces the sample rate, performs partial channel selection filtering, and converts the signal to IQ format.

IQ Receiver

A dual ADC is often used in receiver applications as an IQ pair. Silanna's parts have been designed with that in mind paying attention to good matching between the two ADCs. Despite these efforts, but primarily due to IQ mismatch originating from the circuitry prior to the ADCs, it is often necessary to adjust the offset, gain, and phase of the I and Q signals digitally to compensate for the mismatch.

Silanna's feature-rich dual ADC parts provide digital adders and multipliers for applying offset and gain correction to IQ signal. Phase error correction uses a block that adds or subtracts small amount of Q signal from I and vice versa as depicted in Figure 15. The part does not determine the correction values automatically. The user needs to write those in the DSP control registers using SPI. Decimation by two or four and bypassing the decimator are supported in this mode.

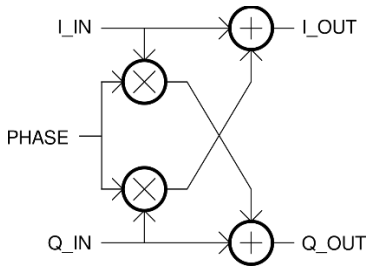

Figure 15. IQ-phase correction block.

The parts also provide a digital mixer (Figure 16) that can be used to frequency shift the signal prior to the decimator. This feature can be used to bring a low-IF signal to base band or to center a signal channel on DC withing a digitized signal band.

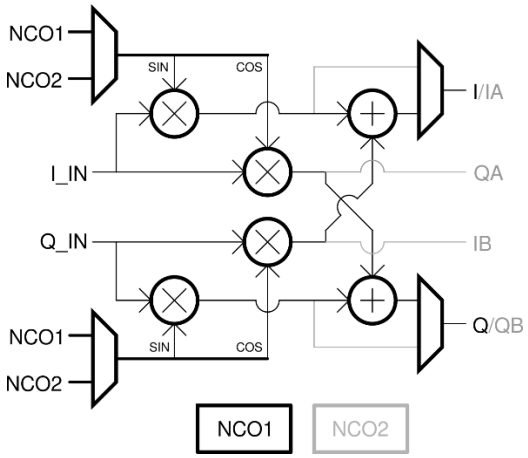

Figure 16. DDC configured for IQ receiver.

Controlling the NCO

Silanna's DSP implements the NCO using a direct digital synthesizer (DDS). It consists of a 48-bit phase accumulator and 16-bit amplitude lookup table. Both sine and cosine outputs are provided.

The phase accumulator, which produces a ramp signal representing the phase range from 0 to 2π, equivalent to one cycle of a sine wave, is incremented every clock cycle. The phase increment, which is related to the DDS output frequency, can be calculated as $\Delta\varphi = \frac{f}{F_S} * 2^{48}$, where f is the output frequency. The smallest possible phase increment (1 LSB) defines the DDS frequency resolution as $\Delta f = F_S / 2^{48}$. Negative frequencies can be obtained by locating the DDS output frequency in the second Nyquist zone, which is accomplished by calculating the frequency as $f \% F_S$, where % is the modulo operator.

The 48-bit frequency control word is spread over three control registers. To guarantee simultaneous application of all the bits, the new frequency is first written to the three registers and then a load bit is toggled in a fourth register.

A 16-bit phase offset can be added to the phase accumulator output. It affects both cosine sine components and it is activated with its own load bit in the same manner as the frequency control word.

Notes

The user should be aware that the latency of the part is not deterministic when the decimator is enabled. The latency of the two ADCs always matches but can vary by one (in the case of decimation by two) or up to three (in the case of decimation by four) sampling clock periods from power cycle to power cycle. This may impact the use of the part in applications where multiple ADC chips are used. The part does not provide a mechanism for synchronizing the start phases of NCOs across multiple chips either.

Revision History

Version	*Date*	*Description of Change*	*Author*
0.1	10/14/2025	Initial Release.	Mikko Waltari
1.0	10/22/2025	Fixed typos and improved the language. Doc Control Initial Release	Mikko Waltari