



HUNT ENGINEERING
Chestnut Court, Burton Row,
Brent Knoll, Somerset, TA9 4BP, UK
Tel: (+44) (0)1278 760188,
Fax: (+44) (0)1278 760199,
Email: sales@hunteng.co.uk
<http://www.hunteng.co.uk>
<http://www.hunt-dsp.com>



The Theory of Digital Down Conversion

Rev 1.2 T.Hollis/R.Weir 26-06-2003

Introduction

The advent of larger and faster Xilinx FPGA's has opened up the field of digital signal processing. The large array of configurable logic blocks within the FPGA give great flexibility together with speed, once configured the FPGA is not as flexible as a processor but is much faster. For many DSP applications speed is important especially for the initial processing of the data, after which the data rate reduces and becomes more manageable. Digital Down Conversion is a technique that takes a band limited high sample rate digitised signal, mixes the signal to a lower frequency and reduces the sample rate while retaining all the information.

The main advantage of using an FPGA to implement the Digital Down Converter is the speed, but it also has advantages associated with any digital signal processing system in that once it is defined it is fixed relative to the sample frequency, and will not change with time or temperature. For example once the coefficients for a digital filter have been set the characteristics of that filter are defined, if an identical filter is required then implement the same length filter with the same coefficients. Identical filters are particularly useful in quadrature sampling systems.

The use of Field Programmable Gate Arrays means that while developing your system if, for example, the filter characteristics are not quite right all that is required is to reconfigure the FPGA with a different filter. Dedicated hardware would most likely require component changes.

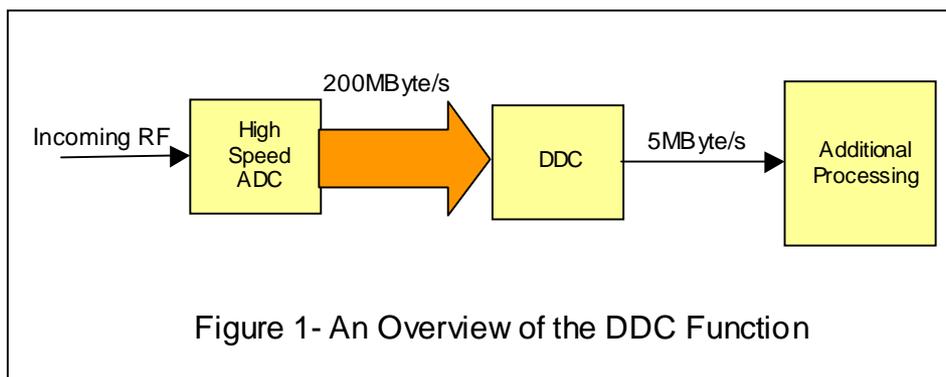
HUNT ENGINEERING manufacture FPGA modules together with A/D and D/A converters on them with sample rates of 100MSamples per Second, which makes them suitable for digital radio applications that would require Digital Down Conversion.

The Digital Down Converter or DDC

A fundamental part of many communications systems is Digital Down Conversion (DDC). Digital radio receivers often have fast ADC converters to digitise the band limited RF or IF signal generating high data rates; but in many cases, the signal of interest represents a small proportion of that bandwidth. To extract the band of interest at this high sample rate would require a prohibitively large filter. A DDC allows the frequency band of interest to be moved down the spectrum so the sample rate can be reduced, filter requirements and further processing on the signal of interest become more easily realisable.

Consider a radio signal lying in the range 39-40MHz. The signal bandwidth is 1MHz. However, it is often digitized with a sampling rate over 100MSamples per Second, representing in the region of 200Mbyte/second.

The DDC allows us to select the 39-40MHz band, and to shift its frequency down to baseband and in doing so reduce the sample rate, with a 1MHz bandwidth, a sample rate of 2.5MHz would be fine - giving a data rate of only 5Mbyte/second. This is shown in Figure 1.



How It Works

A Digital Down Converter is basically complex mixer, shifting the frequency band of interest to baseband. Consider the spectrum of the original continuous analogue signal prior to digitisation, as shown in Figure 2, because it is a real signal it has both positive and negative frequency components. If this signal is sampled by a single A/D converter at a rate that is greater than twice the highest frequency the resulting spectrum is as shown in Figure 3. The continuous analogue spectrum repeated around all of the sample frequency spectral lines.

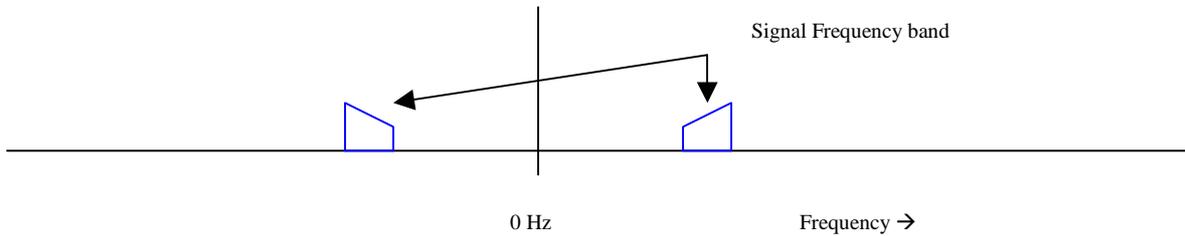


Figure 2 – Spectrum of a continuous analogue signal

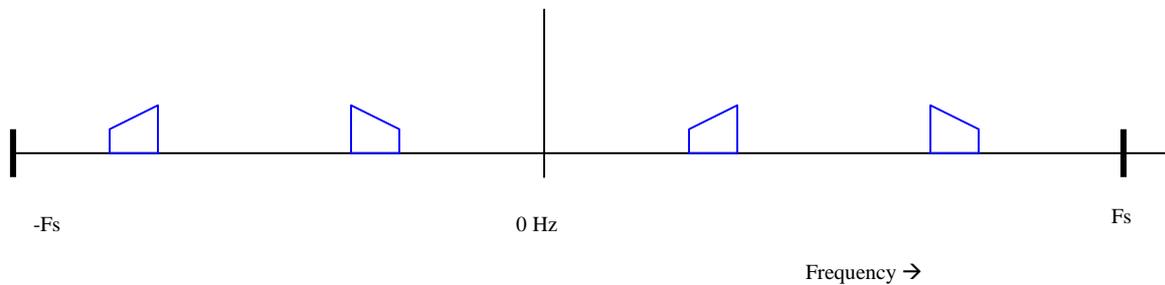


Figure 3 – Spectrum of digitised signal
Sample frequency F_s

The first stage of the DDC is to mix, or multiply, this digitised stream of samples with a digitised cosine for the phase channel and a digitised sine for the quadrature channel and so generating the sum and difference frequency components. Figure 4 shows the amplitude spectrum of either the phase or quadrature channel after mixing, the mixer frequency has been chosen in this example to move the signal frequency band down to baseband. The amplitude spectrum of both phase and quadrature channels will be the same but the phase relationship of the spectral components is different. This phase relationship must be retained, which is why all the filters in the phase path must be identical to those in the quadrature path. It should also be noted that because we have quadrature signals the spectral components from both positive and negative frequencies can be overlaid, for non quadrature sampling the two frequency components would have to be kept separate and so requiring twice the bandwidth.

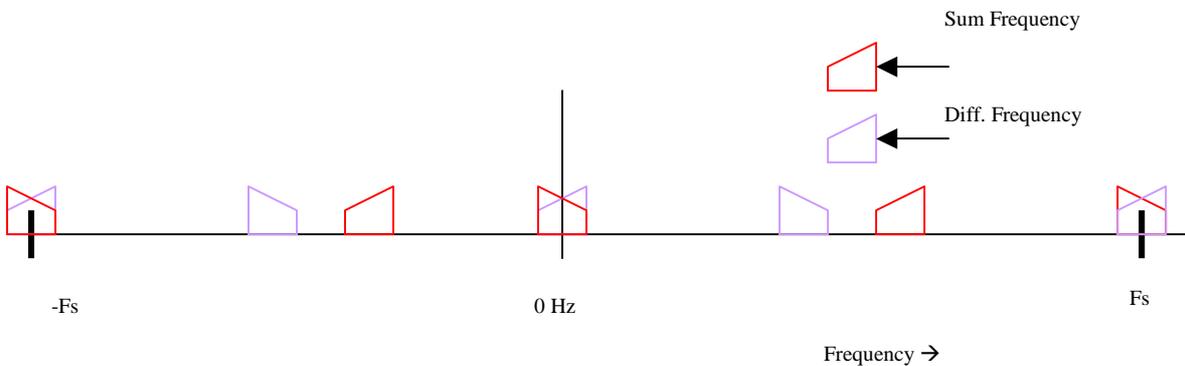


Figure 4 – Spectrum of digitised signal after mixing
Sample frequency F_s

This spectrum of both phase and quadrature signals can now be filtered using identical digital filters, with a response shown in Figure 6, to remove the unwanted frequency components. The phase and quadrature samples must be filtered with identical filters, which with digital filters in an FPGA is not a problem. A digital filter frequency response is always symmetrical about $0.5F_s$. The unwanted frequency components fall outside the pass bands of the filter, giving the resultant spectrum for both phase and quadrature as shown in Figure 7.

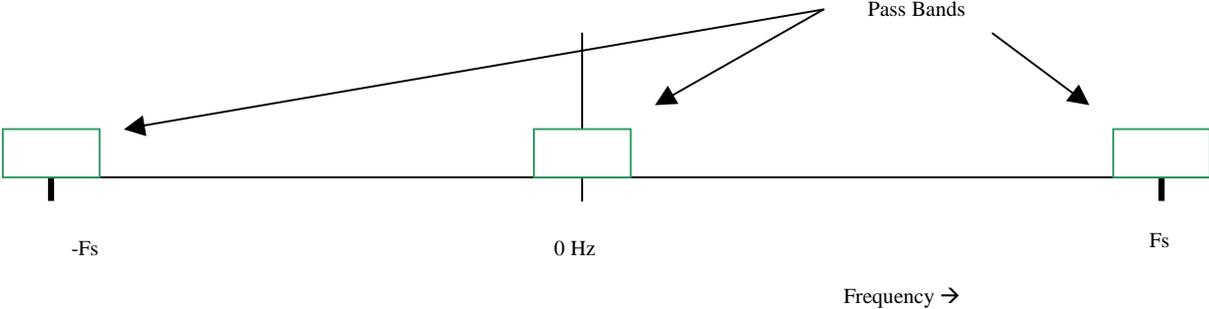


Figure 6 – Low pass digital filter frequency response
Sample frequency F_s

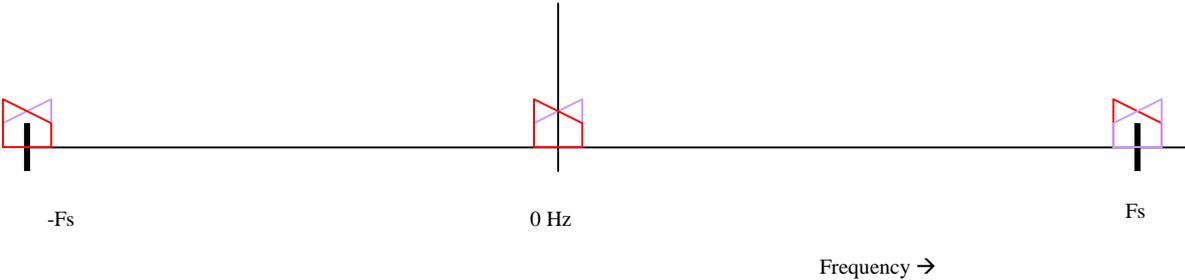


Figure 7 – Spectrum of digitised signal after filtering
Sample frequency F_s

The sample frequency is now much higher than required for the maximum frequency in our frequency band and so the sample frequency can be reduced or decimated, without any loss of information.

Up to this point we have talked about using the requirements for the quadrature signals in the digital mixer and in the filters on the phase and quadrature paths, but now we can start to see the advantages of using this technique. The phase relationship between the phase and quadrature signals when combined as a complex value allow the components from the negative frequency signals to be cancelled. This means that the spectrum of the combined phase and quadrature signals does not have the positive and negative frequency components overlaid as shown in Figure 7, but has just the required positive signal component. In theory the sample frequency of the quadrature signals could be reduced to just greater than the bandwidth of our frequency band, although in practice a slightly higher frequency would be chosen to allow a guard band. The new quadrature spectrum is shown in Figure 8. with the sample frequency reduced to a seventh of the original value. This decimation is achieved by only taking every seventh value and ignoring the rest. This technique of decimation is limited to integer decimation rates, but there are other more complicated techniques that can give non integer decimation, this is covered in more detail later.

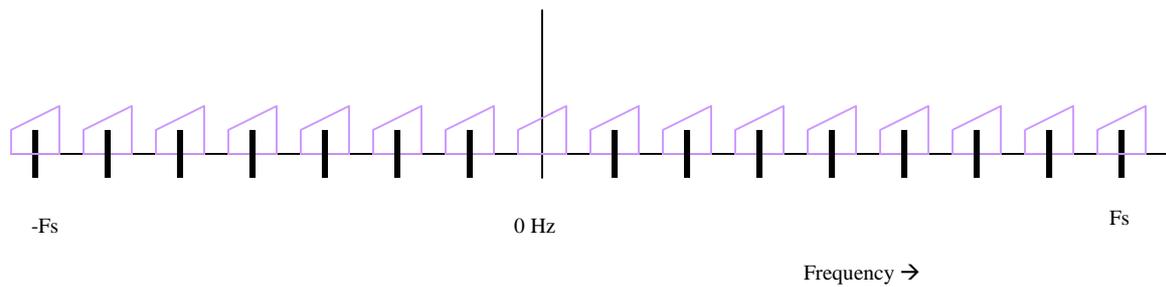


Figure 8 – Spectrum of quadrature signal after decimation
Sample frequency $F_s/7$

The output from the Digital Down converter has retained all the information in our frequency band of interest but has moved it down to baseband and so allowed the sample frequency to be greatly reduced. This has the advantage of simplifying any further processing on the data, together with the gains of possibly time shared functions within the FPGA due to the lower clock frequency, so allowing more processing to be fitted into the FPGA, also the lower effective clock frequency will reduce the power requirement for the FPGA.

The above example is to demonstrate the basic operation of a Digital Down Converter but the detail of a particular converter will depend on the application. The maximum A/D sample frequency may not be high enough so that band pass sampling techniques, as discussed in more detail later, may be used. The signal in the example has been mixed down to base band to demonstrate the maximum reduction in output sample frequency, but this may not be suitable for your application. The frequency shift is determined by the quadrature oscillator frequency.

The filter design is a function of the application requirement, broadband, narrowband, linear phase?

Why Linear Phase Filters?

We mentioned that the filters used to reduce the bandwidth of the signal usually have linear phase characteristics. Linear phase filters are usually more complex than those with arbitrary phase characteristics, so once again, there is a good reason for this.

Communications systems often depend on the relationships between multiple carriers. These carriers may be the same frequency but with different phase; or they may be completely different frequencies. In either case, disturbing the phase relationships would be a bad thing.

For this reason, most DDC designers will try to use linear phase filters exclusively. These appear as a simple delay to the signal, and as all elements of the signal are delayed by the same amount, the signal's integrity is preserved.

How do we decimate by non-integer ratios?

In our description, we described decimation as simply “throwing away” samples. This is a valid thing to do for integer changes in the sampling rate; for example, when decimating a bandwidth limited signal from 3MHz to 1MHz it is perfectly OK to discard every second and third samples.

However, this is not always the case. In some systems we will want to decimate by non-integer ratios – for example, 2,5MHz down to 1MHz. This represents a decimation ratio of 2.5, or 5/2.

In this case, we first INCREASE the sampling rate. This is known as interpolation. We add samples, then filter the signal back to its original bandwidth. Here, we double the sample rate up to 5MHz.

Once the sample rate has been increased, we can once again decimate by discarding samples. In the example, we discard four samples from every five, giving us the 1MHz output rate.

All these concepts are drawn together in the block diagram of Figure 3

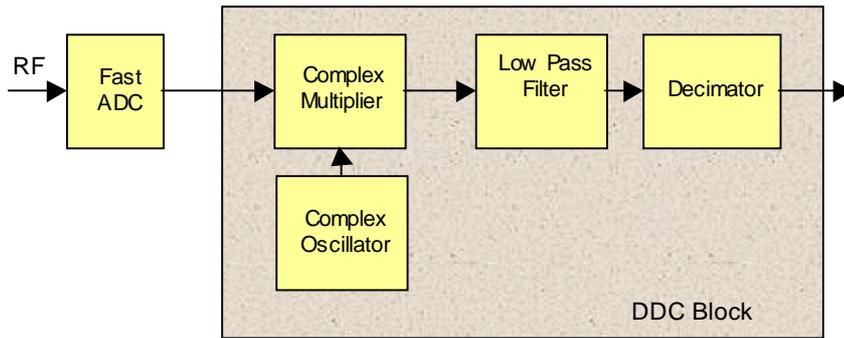


Figure 3 - A Theoretical DDC Block Diagram

How fast an ADC do I need?

The concept of the DDC is to sample the whole input signal and to use digital techniques to reduce the data. However, that may require an unrealistically fast ADC.

For example, Nyquist's theory states that we should sample at a rate "at least double the bandwidth of interest". If we do this on a 1GHz carrier, we would need an ADC sampling at well over 2GHz; that ADC would be protected by anti-aliasing filters, removing any signal above 1GHz. However, this is beyond what can be achieved with today's technology.

Frequency shifting is required before the ADC. This can be done in analog, using an IF stage; or it may be possible to use a composite approach.

Typically, if the signal is at 1GHz, the bandwidth of interest may be only a MHz wide. The filters required to select such a narrow band would be large and complex; but selecting a band perhaps 40MHz wide could be practical. If we can select this band using analog filters, it is possible to sample at a reduced rate – using the bandpass sampling mentioned earlier. In this case, a sample rate (F_s) of 100MHz would mean that our signal could be frequency shifted by $10F_s$, bringing the 1GHz signal in at a much more reasonable 100MHz sample rate.

The more aggressively this technique is applied, the greater the strains on the anti-aliasing filters. It also places great emphasis on high analog bandwidth for the ADC, and extremely low jitter for the ADC clocks; any errors here will be magnified greatly. However, it can be used very effectively.

This technique is complementary to frequency-shifting the signal to an "Intermediate Frequency" or IF. Typically the approach used – direct conversion, IF or "bandpass sampling" will be chosen dependent on the frequencies involved and other system issues –even down to the type of antenna deployed. Most systems use a combination approach, using an analog IF stage to bring the signal down to something that can then be processed digitally.

Why would we use a DDC over Analogue Techniques?

Plainly a DDC is implementing something which could be done in analogue – it's sometimes good to stop and check why we'd want to do this.

The DDC is typically used to convert an RF signal down to baseband. It does this by digitising at a high sample rate, and then using purely digital techniques to perform the data reduction.

Being digital gives many advantages, including:

- Digital stability – not affected by temperature or manufacturing processes. With a DDC, if the system operates at all, it works perfectly – there's never any tuning or component tolerance to worry about.
- Controllability – all aspects of the DDC are controlled from software. The local oscillator can change frequency very rapidly indeed – in many cases a frequency change can take place on the next sample. Additionally, that frequency hop can be large – there is no settling time for the oscillator.
- Size. A single ADC can feed many DDCs, a boon for multi-carrier applications. A single DDC can be implemented in part of an FPGA device, so multiple channels can be implemented – or additional circuitry could also be added.

However, there are some disadvantages:

- ADC speeds are limited. It is not possible today to digitise high-frequency carriers directly. There are techniques to extend the range of ADCs, but often it is simpler to use analogue circuits to bring the carrier down to an IF that digital circuits can then manage.
- ADC dynamic range is limited. In many communications systems, the signal's amplitude can vary greatly. Fast ADCs often only have 12bits of resolution – giving an absolute maximum dynamic range of 72dB. It is often better to use analogue circuits in conjunction with the ADC to implement AGC functions to ensure that this range is best used.

In time, more and more systems will use predominantly digital technology. However, the high speeds of many communication systems will ensure that a hybrid approach, using analogue and digital, will be the best route for many systems for a long time to come. The quest for more spectral space will ensure that new systems will use ever higher frequencies, ensuring that analog approaches will be around for a long time to come!

Further Reading

There are many aspects of DDCs and the HERON system that we have touched on lightly in this tutorial. Try the following for more information:

CIC Filters	The original paper on CIC filters is a mathematical work of art – however, that doesn't make it useful for engineering purposes! A lighter introduction is available at: DSPGuru http://www.dspguru.com/info/tutor/cic.htm This references Hogenauer's original paper, but it's not strictly necessary to understand all the maths to get a CIC working!
Complex Maths	Another useful tutorial on the DSPguru website, complex maths explained in simple terminology – essential for a full understanding of much of communications theory: DSPGuru http://www.dspguru.com/info/tutor/quadsig.htm
DDC Chips	Take a look at how the fixed-function DDC chips are structured. Many of these are overkill – they have to be as they need to match the needs of many users, while your design can be application specific. You can still learn from them; Graychip in particular have excellent datasheets: Graychip has recently been taken over by Texas Instruments but you can still find their datasheets at the TI website analog and mixed signal section