FLEXIBLE ARCHITECTURES FOR WIDEBAND SDR CHANNELISATION

John Lillington (RF Engines, Newport, Isle of Wight, UK; john.lillington@rfel.com)
Steve Matthews (RF Engines, Newport, Isle of Wight, UK; steve.matthews@rfel.com)

ABSTRACT

The purpose of this paper is to compare competing techniques for wideband channelisation, and to assess the flexibility of each of these methods in the context of a software defined radio (SDR) receiver. Distinction is drawn between architectures where all channels are equally spaced and of equal bandwidth, and those architectures which afford greater flexibility. Consideration is also given to the practicalities of channeliser reconfiguration, and the ability for channels to operate independently of one another. Particular emphasis is given to analyzing those architectures that are capable of dealing with wide input bandwidths (hundreds of MHz or more) and large numbers of channels.

1. INTRODUCTION

In many existing digital radio receivers, one of the most expensive components is the analogue circuitry required to carry out the initial downconversion prior to digitisation. With current ADC technology providing high-resolution digitisation at sample rates of up to several hundred MHz, it is becoming increasingly attractive to implement the latter downconversion stages digitally. This is most especially the case when a large number of signals are required to be monitored or downconverted at the same time.

The most common technique involves the use of a Digital Downconverter (DDC). The DDC process is often carried out using custom ASIC chips, of which there are many different varieties available, although DDC FPGA cores are also available from many vendors. Typical DDC functionality is illustrated in
Figure 1. Functions are:

1) A frequency shift of -f to centre the required channel at DC, including conversion from real to complex;
2) Filtering to remove all the unwanted out-of-band signal components that would otherwise alias into the passband on decimation; and decimation by a user-specified factor D; this is achieved using a decimating CIC [1]
3) A further decimate-by-4 lowpass filter to correct for the CIC filter shape and apply a user-defined filter to the output.

Usually, in a digital receiver where the DDC output is being fed into a demodulator, the required output rate is close to an integer multiple, R, of the symbol rate where R is typically in the order of 2 to 4.

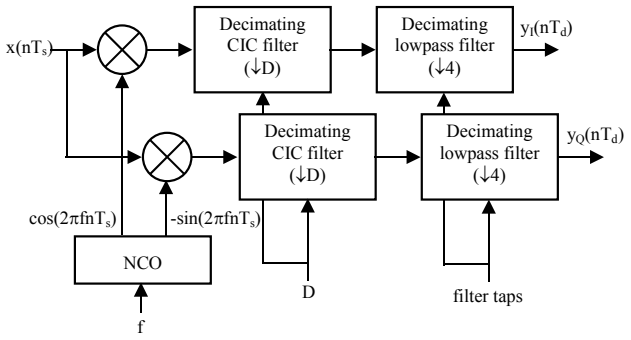

Figure 1: DDC chip architecture

Most DDCs permit the user to control the decimation rate and the filter characteristics, although the precise level of flexibility depends upon the implementation. Typically, a fully programmable DDC ASIC chip supports around 4 independent channels, extracted from digitised inputs, although some reduced functionality DDC chips are now available that support larger numbers of channels (e.g. only supporting a single communication standard). DDC cores for FPGA often have more flexibility, though they can require a lot of silicon.

2. WIDEBAND DOWNCONVERTER SOLUTIONS

There are many cases where several hundred relatively narrowband (100s of kHz) channels are required to be downconverted from a single wideband (100s of MHz) digitised data stream. In this case, it is normally attractive to replace a large number of ASIC DDC chips with a single integrated channeliser.

If the required signal parameters are known at design time, it is possible to eliminate some of the flexibility of the DDC approach to provide far more silicon-efficient Downconverter structures that are tailored to meet specific requirements. Generally, these structures are implemented in FPGA since they are not required for volume applications; however, there is no limitation to their implementation on ASIC. A further advantage of using FPGA technology is that it provides a degree of future-proofing, in that if a different channel structure is required

at a later date, a different FPGA image may be provided that meets this requirement without the need for a complete redesign of the board.

RF Engines have various patented and proprietary channeliser architectures that can be used to meet a wide range of requirements. These may be classified as:

1) Wideband DDC cores providing downconversion of a few relatively wideband sources from a wideband input source.
2) Flexible multichannel Downconverter cores providing channelisation of a large number (a few hundred) of relatively narrowband sources from a wideband input. These channeliser cores can be used to efficiently extract signals from any dynamically selectable frequency with a very wide variety of channel sample rates and filter characteristics.
3) Fixed multichannel Downconverter cores that channelise a very large number (more than a thousand) of channels from fixed channel locations, where the channels have a fixed spacing and all share the same filter shape and output sample rate.

Each of these variants has advantages for different applications. Figure 2 summarises the applications for the different Downconverter techniques. In this figure, wideband DDC cores are seen to provide significant flexibility for a limited number of channels; the fixed Downconverter cores are at the other end of the scale, providing limited flexibility but a large channel capacity. The flexible Downconverter cores occupy the parameter space between the above two where a core can be provided that meets the required flexibility versus efficiency trade-off for a particular application. In general, silicon usage increases both with the number of channels and the required flexibility.

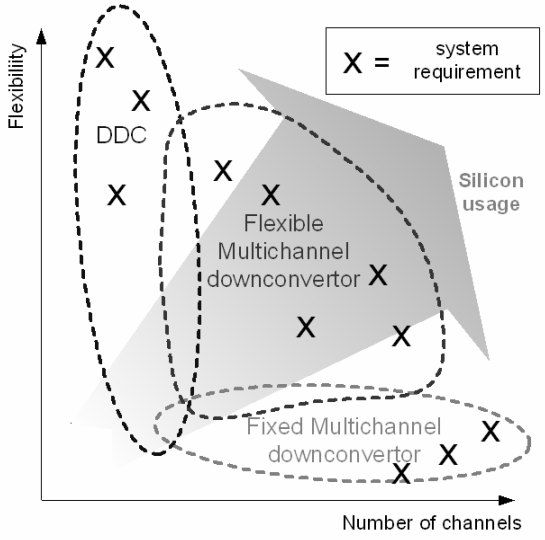

Figure 2: Suitability of Different Channeliser Architectures

The remainder of this paper will concentrate firstly on the implementation of fixed multi-channel downconverters, including mixed radix FFT-based architectures and secondly on flexible multichannel downconverter architectures.

3. FIXED MULTICHANNEL DOWNCONVERTERS

3.1. The WOLA or Polyphase FFT

It is well known that a K-point FFT may be considered as a critically decimating filterbank, providing K equally spaced channels, all filtered by a K-point moving average filter response and decimated by a factor $D = K$. For the WOLA (Weight Overlap Add) FFT, an additional filtering stage is placed prior to the FFT to modify the filter response and to change the decimation factor. The Polyphase FFT technique is similar but permits less flexibility in the selection of the decimation factor. In general, this article refers to the WOLA FFT, though the appropriate implementation should be chosen for each case.

The four design parameters for the WOLA FFT are:

1) The input sample rate, f_s;
2) The length (number of points) of the FFT, K, which provides the channel spacing from the equation, $f_\Delta = f_s / K$;
3) The decimation factor through the WOLA, D, which provides the output sample rate per channel via the relation, $f_{demod} = f_s / D$.
4) The filter impulse response, $\{h[n], 0 \le n < L\text{-}1\}$.

The channel spacing, f_Δ, is fixed by the communication standard and the required sample rate f_{demod} is specified for the demodulator; f_{demod} is normally very close to an integer times the symbol rate.
It is easily shown that for a filterbank to meet the above requirement, integer values of K and D must be found that satisfy:

$$K / D = f_{demod} / f_\Delta \qquad (1)$$

One structure for implementing the WOLA DFT is as shown in Figure 3, where all lines represent complex data. The input is divided into frames of D samples and passed into a delay line. This is then weighted by the filter impulse response, divided into blocks of K samples and overlapped to pass through the FFT. A final step is required to correct the phase of the outputs.

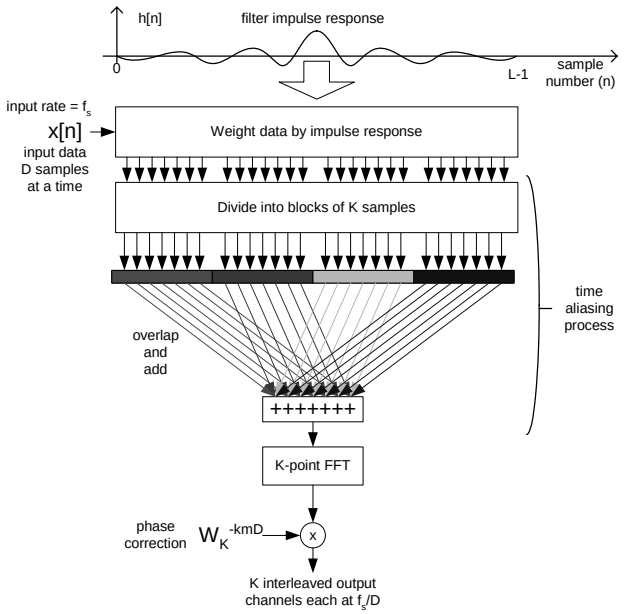

Figure 3: WOLA FFT Structure

3.2. Mixed Radix FFT Designs

The channel spacing, f_Δ, and demodulator rate, f_{demod}, are generally chosen to meet other requirements, resulting in values for K and D that are not powers of 2. For instance, GSM has a channel spacing of 200kHz and a symbol rate of 1625/6 = 270.833kHz. From these numbers and the above equations, a demodulator that requires two samples per symbol provides $D = 48n$ and $K = 65n$, where n is an integer chosen to meet additional requirements such as sample rate, number of channels or usable input bandwidth. For instance, if the requirement was to provide 200 channel outputs, $n = 4$ could be chosen, yielding the solution shown in Figure 4.

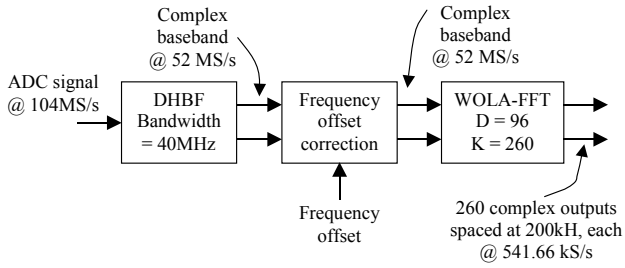

Figure 4: Mixed Radix Solution for 200-Channel GSM Channeliser

In Figure 4, the ADC output is first applied to a DHBF (Distributed Halfband Filter), which converts the real input signal to a complex output at half the rate. Frequency offset correction is then applied to compensate for a systematic frequency offset across all channels, due to Doppler shift or frequency error.

The signal is then passed into the WOLA-FFT, which provides 260 outputs with the required spacing and rate. It is assumed that only the 200 central outputs are required; the remainder are in the analogue anti-aliasing filter transition bands and are discarded.

Alternative values of n could be chosen that would provide different number of channels and FFT length, however, the important thing is that the FFT is always going to have $K = 65n$ points. This means that it must be a mixed radix solution with radices of 13, 5 and factors of n. There is no solution of this type that will provide the required channel spacing and output sample rate without carrying out a DFT of this length - which implies a mixed radix FFT implementation.

RF Engines have generated a design for the above 200 channel GSM channeliser and have shown that, for a 14-bit ADC, it will fit on a Xilinx Virtex-II 6000 FPGA.

In another recent design that required more than 1500 channels to be precisely extracted from a spectrum bandwidth in excess of 40MHz, 2-point, 3-point and 13-point DFT cores were integrated to produce a 1872-point FFT. Again, a radix-2 FFT would not have been able to meet the channel spacing and sample rate requirements for this application. The design fitted comfortably within a Xilinx Virtex Pro50 FPGA.

Although these types of structure form the basis of the majority of channeliser designs provided by RF Engines, they do have one feature in common. They are all restricted to channelisers where all bins have the same filter response and are on a fixed frequency raster. For many applications, this is not a problem but there are some applications where a mix of different channel widths, sample rates and the ability to tune each bin to an arbitrary centre frequency are a requirement. A good example is in satellite communications. The following sections will describe structures which allow this type of flexibility in an efficient manner.

4. FLEXIBLE MULTICHANNEL DOWNCONVERTERS

4.1. The Tunable PFT

The principles behind the Pipelined Frequency Transform (PFT) have been dealt with in several papers such as References [2] and [3]. In its simplest form, a simple Radix-2 PFT achieves its channelisation by a process of frequency band splitting, as shown in Figure 5 below. In this form, however, the silicon efficiency would be very low due to the sample rate reduction at each stage.

A hugely more efficient structure is shown in Figure 6 where, by interleaving the samples at each stage, full use is made of the available silicon sample rate. This patented structure also allows simplification of the complex up / downconversion required at each stage and can, if required, be realized as a multiplier-less architecture. The other key feature of this structure is the *simultaneous*

availability of outputs at each stage of resolution, providing the basis for a flexible *multi-resolution* filter bank.

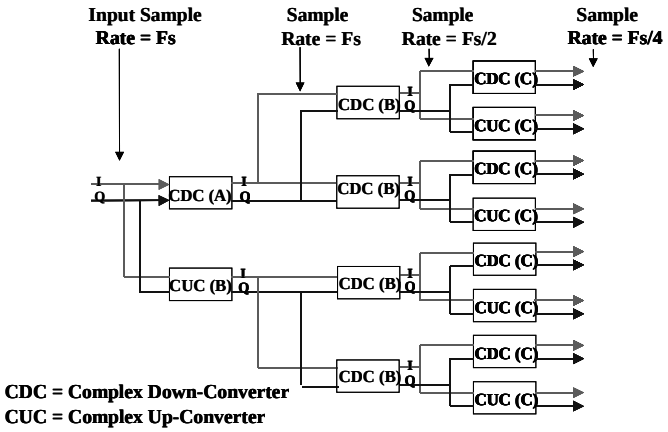

Figure 5: PFT – Simple Tree System

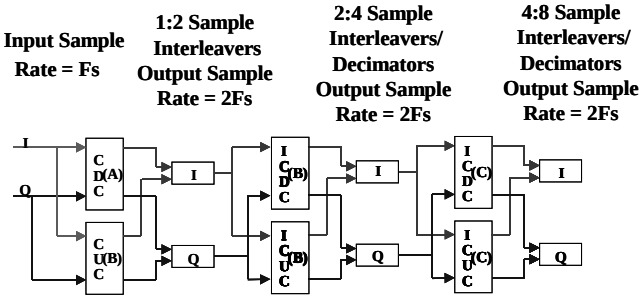

Figure 6: PFT Using Interleaving Architecture

This, in itself, can be a valuable feature and has already been used in various designs. Like the structures examined in Section 3 above, however, although it provides a multi-resolution capability, it still has a set of fixed frequency rasters – i.e. the filters at each stage have a fixed frequency separation and centre frequency and have a common filter design for the bins at each stage.

The Tunable PFT was designed to overcome these restrictions and give a truly flexible architecture. The basic structure is shown in Figure 7 below where the selected bins from each stage are interleaved into a single complex stream. This is possible since, at this point, there is an integer relationship between the sample rates for each stage. Fine tuning of each filter centre frequency may be achieved by passing the interleaved samples through a single polyphase structure consisting of a complex up or downconversion (CUC / CDC) and a final channel filter. The latter allows each channel to have the required filter response (e.g. root-raised cosine). The only remaining requirement, to allow efficient demodulation, is a multi-rate section which allows the final sample rate to be more accurately matched to a multiple of the symbol rate. This is a common requirement for both fixed and flexible downconverters and is dealt with separately in Section 5 below

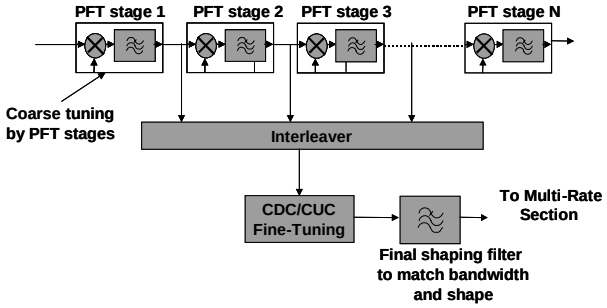

Figure 7. Schematic of Tuneable PFT Architecture

4.2. Hybrid Tunable Structures

It is frequently the case that the widest channel required is still a small fraction of the overall bandwidth being channelised. Since, for the case of a single resolution filter bank, the PFT is likely to be less efficient and have higher latency than, for example, the Polyphase FFT, it can make sense to realise the first part of the flexible channeliser in the latter form. This can then be followed by a tuneable PFT to achieve the flexibility required.

It may also be more efficient, in cases where a wide difference exists between successive stages of filter resolution, to replace some of the final stages with more conventional decimating filters. This could be, for example, a conventional CIC / Decimating FIR or one of the decimating multi-rate structures discussed in Section 5 below.

Recently RF Engines have investigated hybrid architectures that support a maximum aggregate output bandwidth which may be divided between a large number of narrow band channels, relatively fewer wideband channels, or some combination thereof. The user may reconfigure channels at run-time providing the overall bandwidth limit is not exceeded. For example, a design has been shown that will support an aggregate channelisation bandwidth of 32 MHz. This can be composed of up to 16 channels with bandwidths $\leq$ 2 MHz, 32 channels with bandwidths $\leq$ 1MHz, or 64 channels with bandwidths $\leq$ 512 MHz. These techniques exploit resource sharing principles to ensure that silicon resources are minimised for any particular configuration, and show great promise for future flexible channelisation designs, particularly where there is a wide variation in channel sizes.

4.3. Example Tunable Downconverters

Using a selection of the tuneable downconverter techniques described above, RF Engines have shown that a 64-channel flexible channeliser can be generated that will fit comfortably within a Xilinx Virtex II Pro 30 FPGA. This design supports two 16-bit ADC inputs at rates up to

140 MS/s, it provides independent channel tuning with a resolution of 0.01 Hz, and allows decimation factors from 128 to 8192. When combined with a multi-rate filter this architecture is a suitable replacement for typical ASIC or FPGA based DDCs, offering similar flexibility and performance with far greater efficiency.

Many other variants have also been realised, and it has been shown that it is possible to tailor key performance parameters for a particular requirement in order to minimise the FPGA resources. The key parameters that drive the FPGA resource usage have been found to be: the number of downconverter channels; the level of spurious free dynamic range; the minimum level of decimation; and the number of channel filter shapes.

5. MULTI-RATE STRUCTURES

Whilst the architectures described above are able to flexibly filter and downconvert narrow-band channels from a wideband input spectrum, the sample rate of the resulting channelised signals are often determined by fixed decimation factors through the design. Typically, these sample rates are not suitable for the subsequent processing, and hence an efficient additional structure is required to resample each signal to produce the desired sample rate.

RF Engines have demonstrated a highly efficient architecture for this purpose that can resample many channels in an interleaved fashion. Use of fractional resampling techniques allows the channel sample rates to be selected with a resolution which is better than 0.01 Hz.

In one example, the company has shown that 512 channel down-sampler can be implemented in FPGA which supports a maximum aggregate sample rate of 180MS/s. Each channel has separate rate control such that the resulting channel sample rates are fully independent of each other. When targeted at a Xilinx Virtex II Pro 50 the design requires 12% of available logic slices, 7% of multipliers and 16% of block RAMs.

6. SUMMARY

A range of approaches exists for performing the downconversion function in a digital radio receiver, including the classic DDC, FFT based architectures, and novel approaches such as the TPFT. In general there is a trade-off between the level of flexibility offered by the architecture and the silicon resources required for implementation.

The DDC offers excellent flexibility with the user able to select the bandwidth and centre frequency of a channel with high level of resolution. However, this approach suffers from the disadvantage of requiring a large amount of silicon which may preclude its use in systems with more than a few channels.

FFT based approaches, such as the WOLA and Polyphase FFT, are at the opposite corner of the flexibility/resource space. These approaches are highly efficient, with example implementations supporting several thousand channels on one FPGA. However, the inherent use of the FFT requires that all channels must have equal bandwidths, and must be regularly spaced across the input bandwidth.

The TPFT and hybrid variants offer an excellent compromise between these two extremes. A 64-channel downconverter has been shown which offers flexibility which is comparable to a standard DDC architecture and fits comfortably within a Xilinx Virtex II Pro 30. Architectures such as these are a cost effective solution for down-conversion in multi-channel digital receivers, and represent a critical building block for flexible software defined radios of the future.

7. REFERENCES

[1] E. B. Hogenauer. “An economical class of digital filters for decimation and interpolation”, IEEE Transactions on Acoustic, Speech and Signal Processing, ASSP-29(2):155-162, 1981.

[2] PFT Architecture and Comparisons with FFT / Digital Down-Converter Techniques. http://www.rfel.com/download/W02001-PFT White Paper.pdf

[3] J.Lillington. “Comparison of Wideband Channelisation Architectures” International Signal Processing Conference (ISPC), Dallas, 2003