

Comparison of Wideband Channelisation Architectures

John Lillington
CTO, RF Engines Limited
Innovation Centre, St. Cross Business Park
Newport, Isle of Wight, PO30 5WB, UK
Tel: +44 (0)1983 550330
e-mail: john.Lillington@rfel.com

ABSTRACT

The purpose of this paper is to make comparisons between the competing techniques for real-time wideband channelisation including the pipelined FFT, Polyphase DFT's, multiple digital down-converters (DDC's) the Pipelined Frequency Transform (PFT) and its derivative, the Tuneable PFT (TPFT). The main objective is to establish the optimum solution for different application types. Although based on a firm theoretical background, the paper is of a practical nature, giving many real examples which have been designed, placed, routed and tested in FPGA's. Useful information for the system or design engineer includes silicon-performance trade-offs and memory requirements.

General Terms

Algorithms, Architectures, Embedded H/W Design, Radar, SoC/IP Design, Communications, FPGA/CPLD Designs, Military/Aerospace, Sonar, Wireless.

Keywords

IP, Cores, DSP, Wide-band, Real-time, Channel, Pipelined, FPGA, FFT, PFT, Fourier, FIR, CIC, Polyphase, Filter, SoC, Signal, Processing.

1. INTRODUCTION

Whilst high-speed analog to digital converters are available "off-the-shelf" with conversion rates of up to 1.5 Gsps (e.g. Maxim MAX108) and with constantly improving dynamic range, the problems really start with the area of signal processing immediately after the ADC. Typical processing at this stage involves frequency conversion and channelisation and whilst standard digital downconverter (DDC) technologies exist for the selection of narrow-band channels from a medium bandwidth spectrum (e.g. Intersil, TI-Graychip, Analog Devices), they are limited to only a few simultaneous channels for an economical amount of silicon.

Alternative technologies exist where channelisation of the spectrum into equally spaced, equal bandwidth channels is required. These include the FFT (where channel filter performance is not critical) and the Polyphase DFT (where higher performance filters are required). By using pipelining architectures, real-time multi-channel performance can be achieved in a practical amount of silicon. However, real-world situations, such as multi-standard mobile base-stations, software defined radio (SDR) and satellite communications often require channels of non-equal width and spacing and with time varying channel plans. Also, for monitoring, instrumentation and surveillance activities, there is often a need to observe signals in different resolution bandwidths, sometimes simultaneously. For these applications, a novel architecture, the Pipelined Frequency Transform (PFT) and its derivative, the Tuneable PFT (TPFT) are ideally suited since they give maximum flexibility within a practical amount of silicon.

This paper makes comparisons between the main competing techniques for real-time, wideband channelisation including the pipelined FFT, Polyphase DFT, multiple DDC's, PFT and Tuneable PFT. The main objective is to establish the optimum solution for different application types.

2. OVERVIEW of DIFFERENT CHANNELISATION METHODS.

2.1 General Issues.

Before dealing in detail with the various channelisation techniques, a broad overview might be useful. It is not intended to cover some of the more specialist applications such as Chirp-Z transforms or Wavelet filter banks (even though some of the techniques described below may have applicability in these areas). It is those basic techniques which provide multiple channels from a broad band for further processing such as demodulation or signal detection which are of interest here.

It is also worth pointing out that the architectures discussed are generally biased towards hardware implementation such as in FPGA's or ASIC's. Firstly, the processing power required in terms of "multiply /accumulate" operations (MACs) is very high, and, in most cases, way in excess of the peak MAC performance of today's programmable DSPs. Secondly, the most difficult aspect to overcome is the memory bandwidth requirements of a wide band, real-time system. For all the cases considered, it is not clear how this could be achieved without using a totally impractical number of DSPs for the high-end specifications.

2.2 Digital Down-Converters (DDC's).

We will start our overview with conventional digital down-converters (DDC's) These are a well established technique and can be found in COTs products from, for example, Analog

Devices, Intersil and TI-Graychip. It is also relatively straightforward to implement them in FPGA's using custom or standard cores. In cases where only a few channels (typically 4 to 8) need to be selected from the broad band, then such a solution is quite efficient. It also provides a very flexible solution in that each channel can be independently configured for center frequency, bandwidth and filter response. For larger numbers of channels, however, the logic and, more particularly, the memory requirements become excessive, as will be demonstrated later.

2.3 Fast Fourier Transform (FFT).

The Fast Fourier Transform (FFT) and its real-time pipelined implementation are also well established. It provides a very economical solution to the channelisation problem, especially where a large number of channels is required and the channel filter performance is not too critical. It is also generally restricted to cases requiring channels with even frequency spacing and equal filtering.

2.4 WOLA and Polyphase DFT Filter Banks.

An improvement to the filtering performance can be achieved by the use of polyphase filter banks ahead of the FFT, rather than the use of simple "windowing" of the time data. The technique, generally called the "Weight Overlap and Add" or WOLA or its subset the "Polyphase DFT", is becoming more established and is certainly very efficient where large, high quality filter banks are required. Like the FFT, however, it is generally restricted to cases requiring evenly spaced channels with equal filtering.

2.5 Pipelined Frequency Transform (PFT).

A novel form of processing, known as the Pipelined Frequency Transform (PFT) uses a different approach. Based on a "tree" structure, successive splitting and filtering of the frequency band is used to achieve a finer and finer resolution of the broad band. Time interleaving of common processes can lead to a very efficient structure. Advantages include the availability of simultaneous outputs from successive stages, which are at different frequency resolutions and also the ability to independently tailor the filters for different frequency bins. Furthermore, if certain frequency bins or blocks of spectrum are not required, it is simple to exclude them from the processing, leading to greater efficiency.

2.6 Tuneable PFT (TPFT).

In its simplest form, the PFT above still produces equally spaced frequency bins. To overcome this limitation, a derived form, known as the "Tuneable PFT" (TPFT) may be used. This allows independent tuning of the center frequency of all bins as well as independent filters for each bin. Because of the availability of different stage outputs, with different frequency resolutions, the end result is equivalent to having the flexibility of the DDC approach but with the efficiency of the PFT which is important for a larger number of channels.

We will now look at each technique in more detail.

3. DIGITAL DOWN-CONVERTERS.

3.1 General Architecture.

The general architecture of a typical DDC is shown in Fig.1 below.

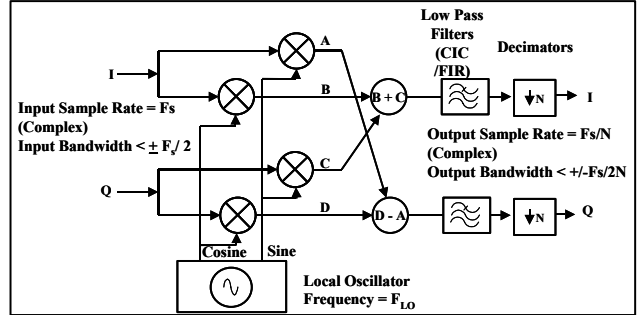


Figure 1 Single Channel Digital Down-Converter using Complex Down-Converter (CDC) and Decimating Filters

Although there are many variants of this design, the principle is broadly the same in each case. The input can be complex or real but, for this discussion, we will assume a complex input. The broad-band input signal is frequency shifted (up or down) to center the required narrow-band channel on zero frequency. This is achieved using a complex local oscillator which is some form of numerically controlled oscillator (NCO) and a complex mixer which, in its basic form, consists of four complex multipliers and two adders. The complexity of the NCO will depend on the final frequency setting accuracy required and on the spurious-free dynamic range (SFDR) of the system.

This is followed by low-pass filtering to extract the required channel, which may consist of any combination of Finite Impulse Response (FIR) filters or Infinite Impulse Response (IIR) filters of which the cascaded-integrator-comb (CIC), Half-Band and decimating FIR filters are typical examples. In the simple analysis presented here, we will assume a multi-stage CIC followed by a decimating FIR. It is usual to perform some of the decimation in the FIR filters, as very high rates of decimation in the CIC require a significant bit growth in the CIC components (see below). Also the FIR filter may need to compensate for roll-off in the CIC passband to achieve adequate passband flatness in the DDC. Typical performance of a single DDC is shown in Fig.2 below.

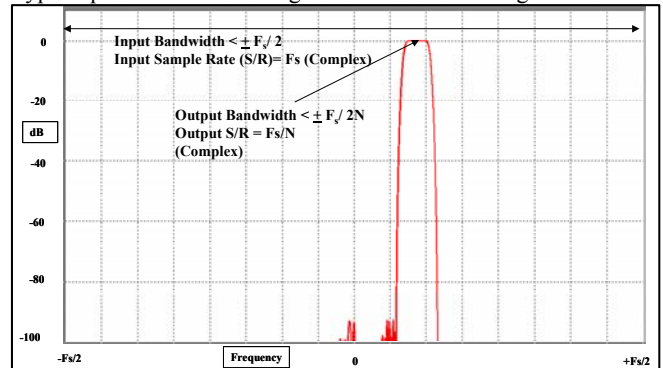


Figure 2 Typical Single Channel Digital Down-Converter (DDC) Frequency Performance.

3.2 Cascaded Integrator-Comb Filters.

CIC filters permit high rate signal decimation (or interpolation) without the need for multipliers and using a very compact architecture (see Ref[1]).

The two basic building blocks of a CIC filter are an integrator and a comb. An integrator (or accumulator) is simply a single-pole Infinite Impulse Response (IIR) filter with unity feedback

coefficient as shown in Fig. 3 below. A typical comb filter is also shown in Fig.3. and behaves as a high-pass filter with a 20dB per decade gain.

Building a CIC filter involves cascading N integrators with N combs, followed by a decimate-by-R block. Although such scheme works, it can be greatly simplified by placing the comb after the decimator. In general, a CIC filter would have N integrator/comb pairs, with N typically ranging from 3 to 6. A 3 stage CIC is schematically illustrated in Fig.3 below.

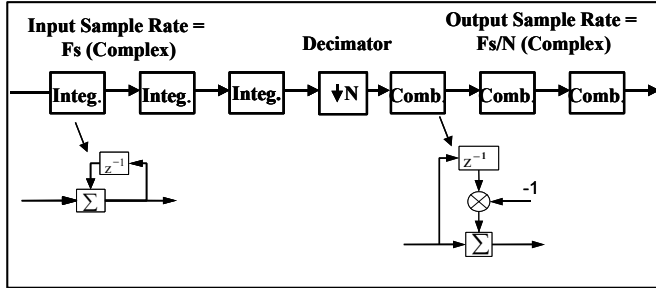


Figure 3 Schematic of a 3-Stage CIC Filter

The magnitude response of a CIC filter can be shown, for large R to approximate to a Sinc function and that spectral nulls will appear at multiples of 1/M Hz. Also, it is clear that the filter attenuation is a function of the number of stages used, while the DC gain is a function of the decimation rate R. When designing a CIC filter, bit-growth must be accounted for, as insufficient bitwidth would lead to an unstable filter.

3.3 DDC Stack Architecture.

It is possible to have a number of DDCs in parallel performing input channelisation. As the number of output channels is equal to the number of DDCs employed in such architecture, it is clear that a linear relationship exists between the amount of silicon and the number of channels required.

It is possible, however, to optimise the DDC stack architecture mentioned above by taking advantage of the changes in sample rate across the system. Since, from each CIC decimator onwards, the comb half of the CICs and the FIRs are clocked at a fraction of the input rate, there is clearly potential for recycling all of the combs and the FIRs into a pipelined version of each. Such an architecture is shown in Ref[2].

4. FFT TECHNIQUES.

4.1 General Architectures.

The whole subject area of FFT techniques is vast with a multitude of algorithms for programmable DSP implementations and a number of COTs ASIC implementations readily available. The intention here is to restrict the discussion to wideband, pipelined hardware solutions, particularly those which are suitable for FPGA realisation. An excellent text on the FFT and its practical realization is given in Ref[3]. Fig.4 below shows a particular implementation of the Pipelined FFT (PFFT) as described in Ref[4]




Figure 4 Basic Pipelined FFT Core

It is based on successive n stages, where 2^n is the size of FFT. Each stage has switched delay elements and butterflies. The switches and delays re-order the data for processing at the next butterfly. There are n butterflies which implement the complex arithmetic, each performs a 2-point DFT and complex phase rotations (twiddles). The input to the first butterfly has a FIFO n/2 buffer stage, which ensures efficient utilisation of the butterfly arithmetic. The normal output of the final stage is bit reversed complex (I/Q) data and thus requires a bit-reverser to achieve normally ordered frequency data.

The choice between Decimate-in-Frequency (DIF) and Decimate-in-Time (DIT) for each butterfly depends on the silicon efficiency of each process and is beyond the scope of this short paper. The reader is referred to Ref[3] for a detailed discussion of this topic.

4.2 Filter Bank Performance.

Fig.5 below shows the effective frequency response of the unweighted FFT, displaying the Sinc/x nature of the sidelobe structure. Comparing this with the filter response of a typical DDC filter (-85 dBc stop-band and low pass-band ripple) shows the clear advantage of the DDC in the filter frequency response.

The standard approach to improving stop-band performance is to weight or “window” the time-domain data. Fig.5 shows, for example, the effect of Kaiser weighting. Whilst the stop-band level is improved, the price paid is a significant widening of the passband. This is demonstrated a little more clearly in Fig.6 below which shows, for example, the equivalent set of overlapping filters for a 32 point complex FFT with a Kaiser window. What this demonstrates clearly is that a signal occupying a narrow frequency band (e.g. CW) will actually appear in a number of adjacent bins in decreasing, but still significant, levels.

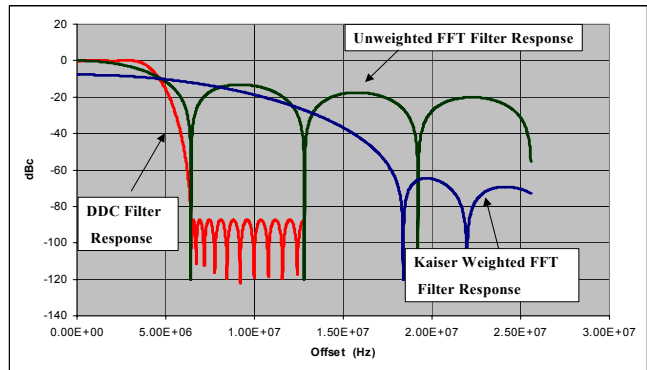


Figure 5 Response of Single FFT Bin – Uniform and Kaiser Windows

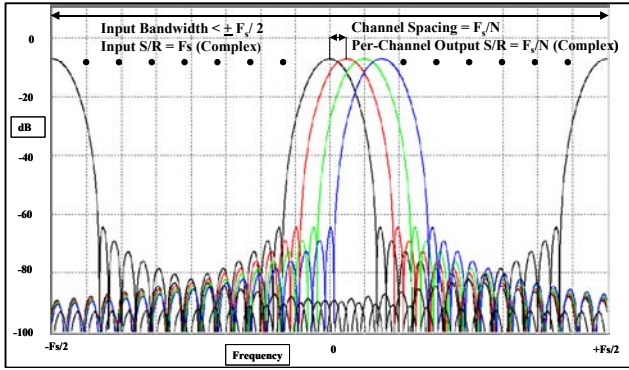


Figure 6 Equivalent Filter Bank Using N-Point FFT (N=32) with Kaiser Weighting

5. WOLA and Polyphase DFT Filter Banks.

5.1 General Architectures.

In the previous section, it was noted that some improvement of filter performance may be achieved by simple windowing of the time domain data. To achieve performance approaching that of a typical DDC (see Fig.5 above) would require a window in the time domain matching that of the overall DDC filter impulse response.




Figure 7 WOLA DFT Filter Bank Implementation.

Thus, for example, a 1024 bin filter bank might require a window some 4K to 5K samples long. This could, of course, be achieved by using an FFT of this length and decimating the output by 4 or 5 but this would be very inefficient, especially for real-time, where parallel processing of several large FFT's would be needed.

Fortunately, there is a much more elegant and efficient solution (See Ref[5]). In its most general form, the "Weight Overlap and Add" or WOLA method is shown in Fig.7 above.

The required filter shape is determined by the weighting function which is L samples long. The weighted data is divided into blocks of K samples to match the DFT length. The blocks are then added together before processing by the DFT. The input data is then shifted along by M samples and the process repeated. In the simple case where M=K then we have a fresh result every K samples so that the system is known as critically sampled (i.e. the

sample rate just satisfies the Nyquist criterion). This may be adequate for some processes (spectral analysis or analysis/synthesis filter bank pairs for example) but often the alias problems caused by critical sampling of a filter bank with finite cut-off rates requires oversampling to some degree. This is simply achieved by making $M < K$ so that the oversampling factor, $I = K/M$ is greater than unity. An advantage of the WOLA is that I need not be an integer. For cases where I can be an integer, a different structure, usually known as the Polyphase DFT, as shown in Fig.8, may be used.

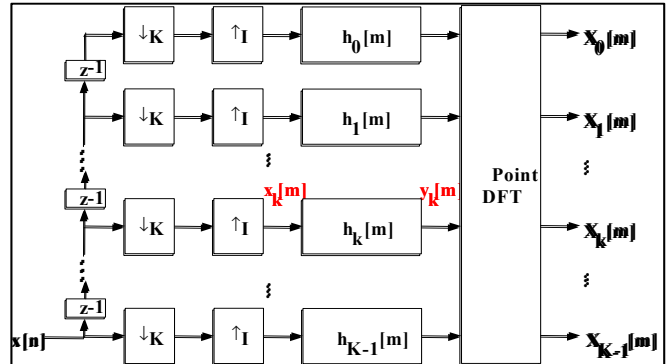


Figure 8 Polyphase DFT Filter Bank Implementation.

For some cases, this may be a more efficient structure, provided that the limitation of integer oversampling is accepted (See Ref[6]). It may also be shown that the Stacked DDC, WOLA and Polyphase DFT give identical results for a given filter response and integer oversampling. The choice between them is mainly based on silicon efficiency.

5.2 Polyphase DFT Filter Bank Performance.

The performance of such a filter bank is illustrated in Fig.9 below. The improvement compared with the standard weighted FFT response (Fig.6) is very obvious.

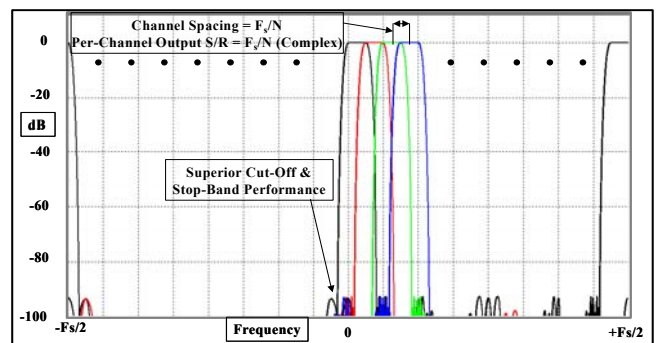


Figure 9 Typical 32-Bin Filter Bank Performance.

6. The Pipelined Frequency Transform (PFT).

6.1 General Architecture.

The underlying concept is one of frequency band splitting where each successive stage of the PFT increases the number of bands by a factor of two, as shown in Fig.10.

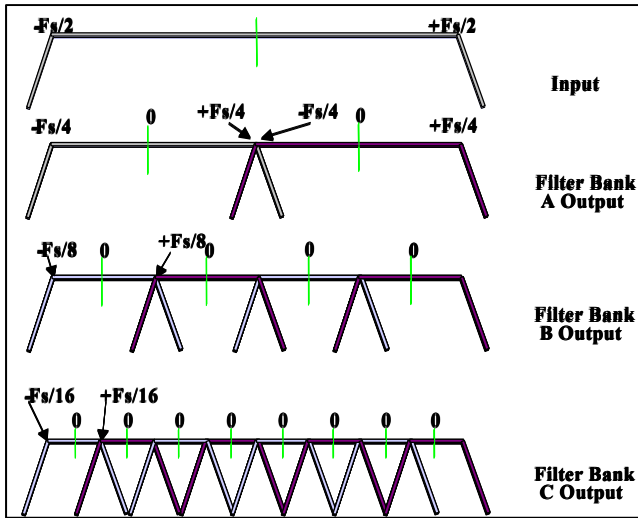


Figure 10 PFT Frequency Band Splitting.

This could be achieved, for example, by a simple tree structure, as shown in Fig.11. The input, which is complex, to preserve positive and negative frequencies, is firstly split into two equal bands using a complex down-converter (CDC) and a complex up-converter (CUC). It would be possible to halve the sample rate for each of the sub-bands since the bandwidth of each has been halved. In practice, a degree of over-sampling is required to avoid image response problems caused by finite filter cut-off rates. Two times over-sampling is used at the output of the first stage. For all successive stages the output is decimated by two, preserving the overall two times over-sampling throughout the system.

The most obvious disadvantage of this approach is that, for large numbers of channels, the tree gets impossibly large. For example, 1024 channels would require 2046 complex (CDC or CUC) modules. Each one of these modules would take the form of Fig.1 which shows, for example, the conventional form of the CDC(A) module, consisting of four multiplies, two adds /subtracts, a sine / cosine look-up table and a pair of low-pass filters. The CUC(A) module would be very similar (differing only in the signs of the adder / subtractor elements). Successive stages would also be very similar except that the local oscillators are now at $F_x/8$ (where F_x is the input sampling rate for the stage) and the output is decimated-by-two.

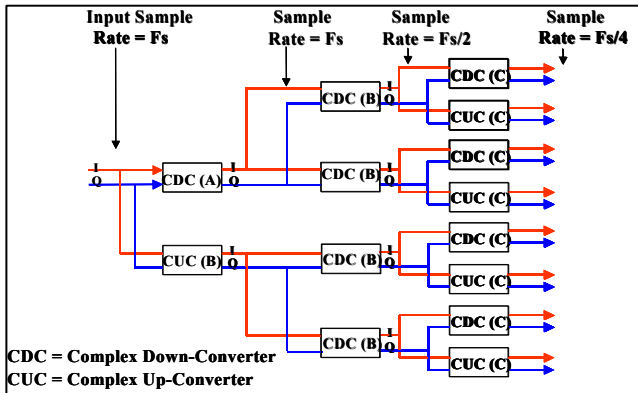


Figure 11 PFT - Simple Tree System.

6.2 Simplification of the Architecture.

Fortunately, the architecture can be greatly simplified in several significant ways. Firstly, with the tree system, the sampling rate drops by a factor of two at each stage. This would lead to inefficient use of the hardware which is capable of running at the full rate, F_s . The most processing-intensive part of each stage lies in the low-pass filters and, since these take an identical form within any given stage, interleaving techniques may be used to regain full efficiency. This involves interleaving the samples for each of the branches within a given stage and modifying the filters (which are normally FIR filters) by adding extra delays between the coefficient multipliers. This is illustrated in Fig.12.

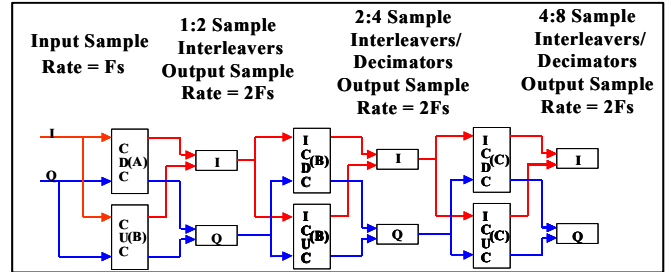


Figure 12 PFT Using Interleaving Architecture.

There are several other simplifications which save silicon including the avoidance of look-up tables and multipliers (see Ref[2] for a more detailed treatment).

6.3 PFT Filter Bank Performance.

The PFT passband performance is very similar to that of the Polyphase DFT, illustrated in Fig.9 above. There are differences, however, in the stop-band, because the PFT is a cascade of filters. This allows higher stop-band attenuation over a large percentage of the broad band, an effect which increases with the number of stages. Also, there are simultaneous outputs available at each stage of the PFT, each giving a different resolution.

Another difference is the ability to use Infinite Impulse Response (IIR) filters at any stage of the PFT, saving silicon for applications where linear phase is not critical and/or where low latency is required.

6.4 The Tuneable PFT (TPFT)

This is an interesting development which makes use of the PFT cascaded structure where intermediate outputs are readily available (see Ref[7]). It is possible, by means of modifying the PFT architecture, not only to extract frequency bands of the desired size, but also to ensure these bands are centered at any given frequency.

This level of *tuneability* is achieved in two stages: firstly the signals are coarsely tuned within the PFT stages, then fine tuned by a complex converter whose Local Oscillator (LO) is a Numerically Controlled Oscillator (NCO) driven by the routing engine. This is shown in Fig.13 below.

The main advantage of performing the tuning operation in two steps is the reduction of size, for a given frequency resolution, of the LUT used for fine tuning. This is because the tuning range required at each successive stage is reduced by a factor of two whereas a DDC would need the fine tuning over the whole input




Figure 13 Schematic of Tuneable PFT Architecture

bandwidth. Overall, the structure is ideal for the replacement of multiple DDC's in applications such as multi-standard base-stations, satellite communications and intelligent antenna systems. A typical scenario is illustrated in Fig.14.




Figure 14 Example Frequency Plan for TPFT

7. SILICON COMPARISONS.

Finally, a brief comparison of silicon usage for the different filter banks will now be made. Within the limited scope of this paper, we will only consider a few examples, based on designs which have been placed and routed in Xilinx FPGA's. Table 1 shows the comparison for filter banks with the following parameters:

Number of bins = 256, 512 or 1024; Filter Stopband=100 dB; Pass-band Ripple=0.1 dB; Filter Overlap=75%; Input Bit-width=14; Sample Rate=102.4e6 Complex, 2x Oversampled

Device: Virtex 2-6000, LUT's=67584, RAM=18432 bits, 18 Bit Multipliers=144

Table 1 Silicon Comparisons for Various Filter Banks

Filter Bank Type	No. of Bins	Bin Spacing (kHz)	Logic (LUT's)	RAM (Bits)	18 Bit Multipliers
Stacked DDC	256	400	317,498	436,224	N/A
	512	200	650,114	876,544	N/A
	1024	100	1,336,754	1,761,280	N/A
Polyphase DFT (Radix 2)	256	400	8,070	4,608	30
	512	200	9,169	4,793	34
	1024	100	10,341	5,345	42
PFT (Radix2)	256	400	27,930	3,840	N/A
	512	200	32,270	6,529	N/A
	1024	100	36,610	10,625	N/A

The most obvious conclusion is that the stacked DDC approach is very inefficient compared with the other two techniques. To be fair, the particular design used did not make use of the dedicated multipliers available in Xilinx Virtex 2 devices. Even so, the use of stacked DDC's for more than about 8 bins is uneconomical.

Direct comparison of the Polyphase DFT and PFT approaches is not simple. The PFT has been configured as a "multiplier-less" design and does not make use of the dedicated multipliers

(although it could do so). Also, it must be remembered that the PFT has outputs available at *each* stage, making it very useful in certain applications. Furthermore, silicon efficiency is much improved if it is only necessary to output bins over selected portions of the broad-band. The general conclusion is that for smaller numbers of bins (up to around 256) the silicon requirements are similar. For larger numbers of bins, the Polyphase DFT gains rapidly, particularly in terms of memory and becomes the preferred choice for single, fixed filter banks

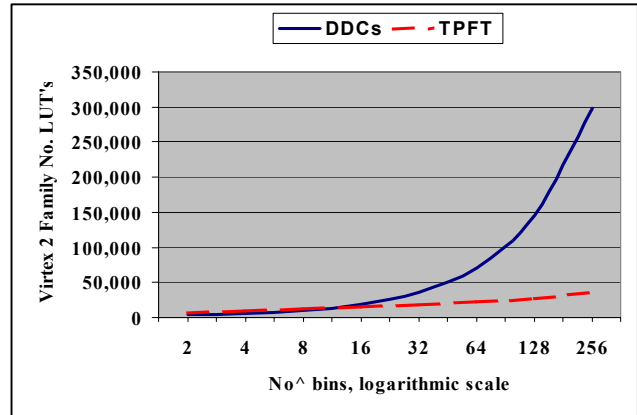


Figure 15 Logic Comparison for DDC's and TPFT

For tuneable filter banks, the best comparison is between stacked DDC's and the Tuneable PFT. Fig.15 compares the logic requirements of the two approaches for up to 256 bins. Above about 16 bins, the TPFT wins rapidly. A similar comparison exists for the memory requirements

8. 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] L.R. Rabiner and B. Gold. "Theory and Application of Digital Signal Processing" Prentice-Hall 1975
- [4] Pipelined FFT <http://www.rfel.com/download/W02004-Pipelined FFT White Paper.pdf>
- [5] Crochiere, RE and Rabiner, LR, *Multirate Digital Signal Processing*, Prentice Hall (Englewood Cliffs, NJ), 1983, ISBN 0-13-605162-6.
- [6] Gumas, CC, "Window-presum FFT achieves high-dynamic range, resolution", Personal Engineering & Instrumentation News, July 1997, pgs 58-64. <http://www.chipcenter.com/dsp/DSP000315F1.html>
- [7] TPFT-Tuneable Pipelined Frequency Transform <http://www.rfel.com/download/W02003-Tuneable PFT White Paper.pdf>

