

WORKLOAD COMPARISON OF DIGITAL DOWN CONVERSION ARCHITECTURES FOR SOFTWARE DEFINED RADIOS

Patrick Crescini (patrick.crescini@navy.mil)¹, Hiba Haidar (hiba.haidar@navy.mil)¹, and fred harris (fred.harris@sdsu.edu)²

¹Space and Naval Warfare Systems Center, Pacific, San Diego, CA, USA

²San Diego State University, San Diego, CA, USA

ABSTRACT

Digital down conversion to base band consumes considerable power in software defined radio receivers. In this paper, we present and compare the workload of standard down converter architectures: one based on a cascaded integrator-comb (CIC) filter, one based on a cascade of half band filters, one based on polyphase down-samplers, and one based on a polyphase filter bank (PFB). Separate comparisons are made for applications requiring down conversion of a single channel and for multiple channels. Workload analysis is expressed as the number of operations per input sample required to deliver output samples due to the sample rate reduction accompanying the bandwidth reduction process. Workload performance for the dual process, digital up conversion is the same as for digital down conversion and thus is not separately performed.

1. INTRODUCTION

A digital down converter (DDC) converts a real digital signal with an arbitrary center frequency to a complex base band digital signal with a center frequency of zero. It typically achieves this with three processes: frequency translation via a quadrature mixer driven by a direct digital synthesizer (DDS), band width reduction via a low pass filter, and sample rate reduction commensurate to the band width reduction [1]. Because of its utility, it is a necessary component in software defined radio (SDR) receivers.

The most simple example of a DDC is the standard single channel down converter shown in Figure 1. Each of its components carries out one of the three typical DDC processes. Of particular concern is the low pass filter, which may be implemented in multiple ways. The most obvious way, shown in Figure 2, is the direct implementation of a finite impulse response (FIR) filter which performs an inner product using multiply and accumulate (MAC) operations. This structure is also known as a tapped delay line.

Suppose we wish to down convert a channel with a bandwidth that is half the output sample rate and a one-sided stop band that begins at 0.75 times the output Nyquist frequency at base band. Also, suppose an out-of-channel attenuation of 100 dB




Figure 1: A standard single channel down converter.

and that the input sample rate is 2,048 times the output sample rate. Then, from [1], we may estimate the length N_{taps} of the direct FIR filter implementation using Equation 1, where f_s is the input sample rate, Δf is the transition band width of the filter (i.e. the difference between the base band down converted channel's one-sided bandwidth and the beginning of its one-sided stop band), and $P_{attenuation}$ is the out-of-channel attenuation. This yields a filter with 37,236 taps, which consists of 37,236 real multiplications and 37,235 real additions per in-phase (I) or quadrature (Q) path.

$$N_{taps} = \frac{f_s}{\Delta f} \frac{P_{attenuation}}{22} \quad (1)$$

We easily see that, even for single channel applications, the tapped delay line FIR filter alone makes the standard single channel down converter computation intensive. The inefficiency of this down converter is further compounded by mixer and filter operation at the input sample rate.

In this paper, we present and compare the workload of three alternatives to the standard single channel down converter. These alternatives are: an architecture based on the cascaded integrator-comb (CIC) filter, a cascade of half band filters, and an architecture based on the polyphase filter bank (PFB). We define workload as the number of operations required to deliver one output sample. Workload analysis of DDCs is significant to minimizing the power dissipation of SDR receivers because it measures economy in operations, which are physically manifested as power dissipating digital circuits.




Figure 2: Direct implementation of an N-tap FIR filter, also known as a tapped delay line.

2. CIC FILTER-BASED ARCHITECTURE

A DDC architecture based on the CIC filter is shown in Figure 3. It is an attractive candidate for a low workload DDC because the CIC filter performs low pass filtering for large integer down sampling without multiplications. As a result, the down sampled CIC filter output is four times the output sample rate. This output is further processed by two half band filters, each followed by 2-to-1 down samplers, to finish channel extraction.


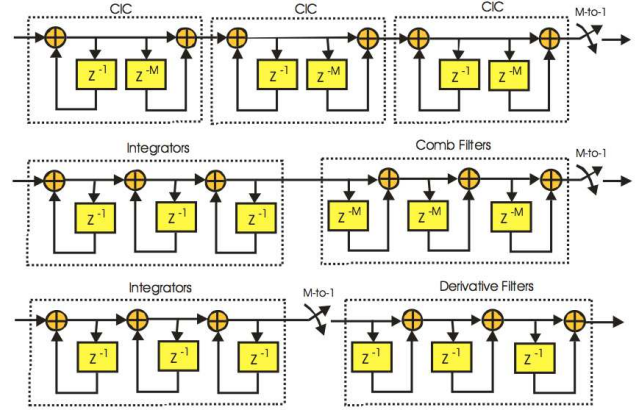


Figure 3: A DDC architecture based on the CIC filter.

2.1. CIC Filter

The single stage CIC filter, which consists of an integrator followed by a comb filter, is a poor filter; its side lobe maximum power level is 13 dB down from that of its main lobe [1]. However, the multistage CIC filter has a side lobe maximum power level that is further down—commensurate to the number of stages—from that of its main lobe. Variations of the multistage CIC filter are shown in Figure 4. Coupled with aliasing, it serves as an efficient filter for large integer down sampling.

The top plot of Figure 5 shows the frequency responses of CIC filters with four, five, and six stages superimposed with the stop band masks. These masks, centered at filter zero-crossings, define regions with a maximum power level of $-P_{attenuation}$ that alias into the pass band following down sampling. From Figure 5, we see that the CIC filter with six stages provides sufficient attenuation at these regions. By examination of Figure 4, we see that each CIC filter stage contributes two additions. Thus, six stages for each I and Q path contributes a total of 24 real additions.


 Figure 4: Variations of CIC filters with embedded M -to-1 down samplers. The bottommost variation is known as a Hogenauer filter.

2.2. First Half Band Filter with CIC Filter Correction

The bottom plot of Figure 5 shows the frequency responses along the pass band of the CIC filters with four, five, and six stages. The responses are mathematically expressed by Equation 2, where M is the down sampling factor serviced by the CIC filter (512) and k is the number of stages (6). They are sinc-like in behavior. As a result, the responses exhibit curvature that distorts pass band signals.

$$H(f) = \left[\frac{\sin(2\pi \frac{f}{f_s} \frac{M}{2})}{\sin(2\pi \frac{f}{f_s} \frac{1}{2})} \right]^k \quad (2)$$

This distortion is corrected by the first half band filter. In addition to band limiting the down sampled CIC filter output for 2-to-1 down sampling, it is designed with the Remez exchange algorithm to have an inverse sinc response along the pass band. Performed in MATLAB, a vector containing equally spaced pairs of closely located normalized frequencies along the pass band is used to generate another vector containing the reciprocal of Equation 3 (the four-term Taylor series expansion of the sinc function) calculated at these frequencies [2]. The resulting vector is provided to the MATLAB `remez` function as part of the

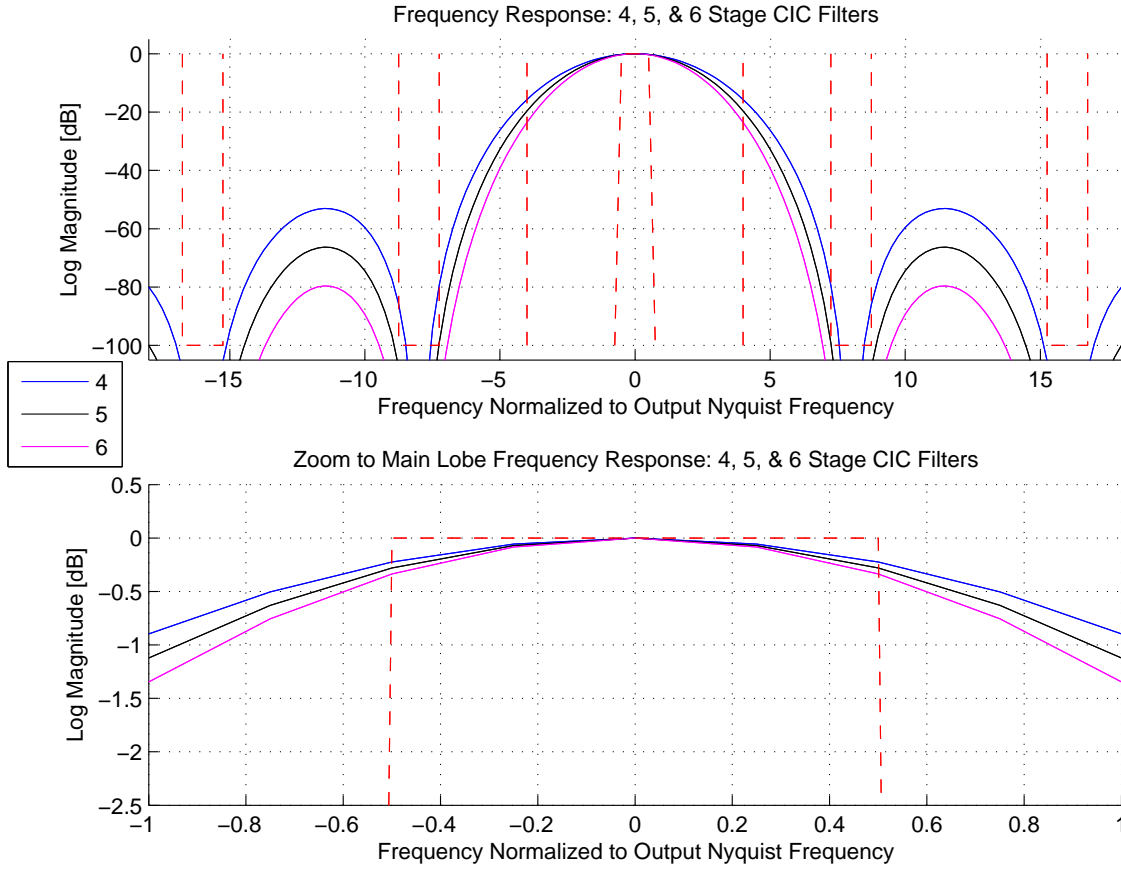


Figure 5: Frequency responses of CIC filters with four, five, and six stages with stop band masks and showing pass band curvature.

filter specifications.

$$\left(\frac{\sin(\theta)}{\theta}\right)^k = \left(1 - \frac{\theta^2}{3!} + \frac{\theta^4}{5!} - \frac{\theta^6}{7!}\right)^k \quad (3)$$

Providing the filter specifications to `remez` as prescribed above yields a 21 tap filter. Using the Noble identity, the 2-to-1 down sampler that follows the filter is relocated to its input in the form of an input commutator. This commutator provides inputs to two 11-tap subfilters $H_0(z)$ and $H_1(z)$ that respectively contain the odd and even coefficients of the designed filter. (The designed filter is zero-packed to contain an even number of taps before division into two subfilters.) Figure 6 illustrates this structure.

The two-path polyphase implementation of this filter contributes 22 multiplications and 21 additions per I or Q path. Thus, the filter contributes a total of 44 real multiplications and 42 real additions. Note that these operations are carried out at the reduced sample rate.




Figure 6: Two-path polyphase filter implementation of the half band filter.

2.3. Second Half Band Filter

The second half band filter band limits the down sampled output of the first half band filter for 2-to-1 down sampling to the output sample rate. Like the first half band filter, it is designed using `remez` and transformed into a two-path polyphase filter shown in Figure 6. Unlike the first half band filter, it does not provide correction to the CIC filter.

The impulse response of the second half band filter is a windowed sinc function. It is given by Equation 4, where $w(n)$ is some windowing function chosen to meet the transition band width and stop band side lobe power level specifications. We see that the zero-indexed coefficient will be 0.5 while the non-

zero even-indexed coefficients will be zero. These coefficients populate the $H_0(z)$ subfilter, while the odd-indexed coefficients will constitute the $H_1(z)$ filter. With respect to number of operations, this half band filter is economic because $H_0(z)$ consists of all delays and with only one arithmetic shift. On the other hand, subfilter $H_1(z)$ contains the coefficients that are not interstitial zeros.

$$h_{LP}(n) = \frac{1}{2} \frac{\sin(\frac{n\pi}{2})}{\frac{n\pi}{2}} w(n) : -N \leq n \leq N \quad (4)$$

Zero-packing a quarter band filter designed with `remez` yields a 57-tap second half band filter. Subfilter $H_0(z)$ contributes 28 additions while subfilter $H_1(z)$ contributes 29 multiplications and 28 additions. Thus, for both I and Q paths, the filter contributes 58 real multiplications and 114 real additions.

2.4. Workload Analysis

To analyze the workload of an architecture, we examine the architecture from its output to its input. For the CIC filter-based architecture to produce one output, the second half band filter requires two outputs from the first half band filter. To generate two outputs from the first half band filter, the first half band filter requires four outputs from the Hogenauer filter. In turn, the Hogenauer filter requires 2,048 inputs to provide four outputs. Table 1 summarizes the workload for a single and multiple (2,048) channel applications.

3. HALF BAND FILTER CASCADE

Yet another popular option to perform the filtering and resampling operation is a cascade of half band filters. Each successive filter in the chain reduces the signal band width by two and reduces the sample rate 2-to-1 so that the next filter operates at a reduced rate. In fact, the 2-to-1 down sampler can be performed at the input to the filter by partitioning the half band filter to be two-path filter. Figure 6 shows this filter structure, which is also used in the half band filters in the CIC filter architecture. The half band filter cascade architecture is shown in Figure 7. More detail about this architecture is provided by [3].




Figure 7: Half band filter cascade DDC architecture.

3.1. Filter Design Methods

The half band cascade consists of 11 filters designed using two different methods: Pascal's triangle and 2-to-1 zero-packing of a `remez`-designed quarter band filter. For the former, Pascal's triangle contains the coefficients of the polynomial $(z + 1)^n$, where z is a number on the unit circle in the complex plane and n is an integer. This polynomial is zero when $z = 1$, which corresponds with the half sample rate. As n increases, the number of zeros at the half sample rate also increase, causing a "widening" of the zero-crossing adjacent to the half sample rate [2]. This zero-crossing accommodates the pass band upon aliasing. The cascade uses filter coefficients in the third and fourth rows of the Pascal's triangle in Figure 8 for the first six filters. Figure 9 provides the impulse and frequency responses of a filter designed with the latter method, which was addressed in the previous section. The seventh through 11th filter were designed with this method.

			1			
		1		1		
	1		2		1	
	1	3		3	1	
	1	4	6	4	1	
1	5	10	10	5	1	
1	6	15	20	15	6	1

Figure 8: Pascal's triangle.

3.2. Workload Analysis

Workload analysis for the half band cascade is similar to that of the CIC filter-based architecture. Table 2 summarizes the workload of the half band cascade.

4. PFB-BASED ARCHITECTURE

Still another option to down convert, filter, and down sample the desired narrow band channel is the polyphase filter partition of a narrow band, band centered filter. This structure is shown in Figure 10. This filter structure is visualized as a low-pass prototype up converted to the input signal's center frequency, applied to the band centered input signal, and then down sampled because it has been bandwidth limited by the filtering process. The down sampling preserves the signal's identity but aliases its spectrum to base band provided that the center frequency is an integer multiple of the output rate. The filter can then be partitioned into an M -path filter with the down sampler at the input rather than at the output. This is akin to the operation performed on the two-path half-band filters. After the polyphase partition, we are able to extract from each path the complex phase rotator

Table 1: Workload Summary of the CIC Filter-Based Architecture

Single Channel						
	Mixer	Integrators	Combs	HB1	HB2	Total
Adds	0	24,576	48	84	114	24,822
Mults	4,096	0	0	44	58	4,198
Total Ops						29,020
2,048 Channels						
	Mixer	Integrators	Combs	HB1	HB2	Total
Adds	0	50,331,648	98,304	172,032	233,472	50,835,456
Mults	8,388,608	0	0	90,112	118,784	8,597,504
Total Ops						59,432,960

Table 2: Workload Summary of the Half Band Cascade

Single Channel													
	Mixer	HB1	HB2	HB 3	HB4	HB5	HB6	HB7	HB8	HB9	HB10	HB11	Total
Taps		3	3	3	4	4	4	9	13	21	21	57	
Adds	0	6,144	3,072	1,536	768	384	192	288	208	168	84	114	12,958
Mults	4,096	0	0	0	512	256	128	160	112	88	44	58	5,454
Total Ops													18,412
2,048 Channels													
	Mixer	HB1	HB2	HB 3	HB4	HB5	HB6	HB7	HB8	HB9	HB10	HB11	Total
Taps		3	3	3	4	4	4	9	13	21	21	57	
Adds	0	12.6M	6.3M	3.1M	1.6M	786K	393K	589K	425K	344K	172K	233K	26,537,984
Mults	8.4M	0	0	0	1.0M	524K	262K	327K	229K	180K	90K	118K	11,169,792
Total Ops													37,707,776

that originally up converted the filter.

There are three significant advantages to this architecture over the more traditional digital down converter. The first is that the signal does not become complex until it leaves the filter as opposed to making it complex via the input heterodyne, on the way into the filter. As a result, redundant components from having separate in-phase (I) and quadrature (Q) components are avoided. The second advantage is that all the filter arms operate at the reduced output rate rather than at the high input rate. This implies that processed samples are not discarded (as in the other architectures), thus wasting energy. The third advantage is that multiple channels can be serviced by this architecture by replacing the phase rotators at path outputs with an IFFT as shown in Figure 11.

4.1. PFB

The subfilters in the PFB obtain their coefficients from a prototype filter. Suppose this prototype filter is the filter derived for the standard single channel down converter in the Introduction. Thus, the prototype filter has 37,236 taps. Because the input sample rate is 2,048 times the output sample rate, we partition the prototype filter into 2,048 paths. In order for the resulting subfilters to have the same number of taps, we must zero-pack

the prototype filter so that it has a length that is divisible by 2,048. This yields 38,912 taps, which divide into 2,048 19-tap subfilters. Therefore, each subfilter contributes 19 real multiplications and 18 real additions.

4.2. Radix-2 Cooley-Tukey IFFT Algorithm

Recall that the IFFT is an efficient implementation of the inverse discrete Fourier transform (IDFT) given by Equation 5. We see that the complexity of computing an IDFT for N points is N^2 . This may be improved with the use of a divide-and-conquer algorithm such as the ubiquitous radix-2 Cooley-Tukey algorithm. This algorithm can be thought of as performing an IFFT on a two-dimensional array containing the lexicographically mapped input sequence. The phase origins of the rows are aligned with phase rotators known as twiddle factors. These twiddle factors contribute multiplications to the IFFT.

$$x(n) = \frac{1}{N} \sum_{k=0}^{N-1} X(k) e^{j \frac{2\pi}{N} kn} \quad (5)$$

The radix-2 Cooley-Tukey algorithm reduces the complexity of computing an IDFT from N^2 to $N \log_2 N$. Because the first two stages of this IFFT consist of multiplications by ± 1 and ro-




Figure 9: Impulse and frequency responses of the ninth filter in the half band filter cascade.

tations about j in the complex plane, it is further reduced to $N \log_2 \frac{N}{4}$. This complexity equal to the number of butterfly computations performed by the IFFT, and each butterfly computation has four real multiplications and six real additions. Because N in our case is 2,048, we have 18,432 butterfly computations. As a result, the radix-2 Cooley-Tukey IFFT contributes 73,728 real multiplications and 110,592 real additions.

4.3. Good-Thomas IFFT Algorithm

The Good-Thomas IFFT algorithm provides additional economy over the radix-2 Cooley-Tukey algorithm by eliminating the need for twiddle factors. By using number theoretical mapping schemes based on the Chinese Remainder Theorem and Ruritanian Correspondence, the two-dimensional input array contains rows that have naturally aligned phase origins. Loading the input sequence into a two dimensional array to be processed is done with a bishop's tour while unloading a processed array is done with a generalized knight's tour [4].

Due to its loading and unloading mapping schemes, this algorithm is valid only for input sequences of composite length with

relatively prime factors. Because 2,048 does not have prime factors, it cannot be used. Rather, an IFFT length of 2,520 with the prime factors of 5, 7, 8, and 9 are used [4]. Due to the fact that the sample rate changed will need to now be 2,520 instead of 2,048, a PFB-based DDC using this IFFT will not be analyzed for workload in this paper.

4.4. Workload Analysis

Analysis of the workload for this architecture is more straightforward for the multichannel case. The analysis for the single channel case can be derived from the multichannel case by analyzing the workload of a single path through the filter bank plus its associated phase rotator. Table 3 gives the workload summary of the multichannel PFB-based DDC.

5. CONCLUSION

For single channel applications, the half band filter cascade is the most practical and efficient architecture. It offers nearly a 1.5:1 advantage over the CIC filter-based architecture with re-

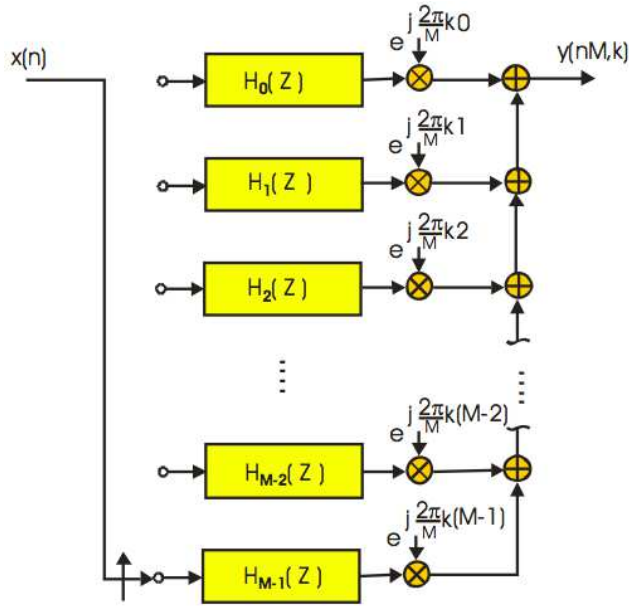

 Figure 10: M -path polyphase partition of a band-centered filter.


Table 3: Workload Summary of the PFB-Based Architecture

	2,048 Channels		
	PFB	Radix-2 Cooley-Tukey IFFT	Total
Adds	36,864	110,592	147,456
Mults	38,912	73,728	112,640
Total Ops			260,096

spect to number of arithmetic operations needed to produce one output sample. Despite the fact that the most prominent filter in the CIC filter-based architecture contains no multiplications, it must perform a substantial number of additions against samples that are eventually discarded. Moreover, for this particular application, the processing overhead contributed by a filter bank in the PFB-based architecture is not immediately necessary in the half band filter cascade.

In the case where multiple channels must be baseband down-converted, the utility of the filter bank in the PFB-based architecture is advantageous and clearly provides the best economy at a two order of magnitude advantage over the other two. This economy is chiefly derived from operation at or near the output sample rate, implying that all processed samples are not discarded. Furthermore, the return on investment in large computational overhead required by the filter bank is evident with the separation of aliases provided by the IFFT.

There exists opportunities to improve the performance of the half band cascade and the PFB-based architecture. For the former, linear phase FIR filter implementation of half band filters can be substituted for recursive filters for additional economy [3]. For the latter, changes in the output sample rate can allow the use of the Good-Thomas IFFT algorithm. Use of this algo-


 Figure 11: M -path polyphase partition of multiple band-centered filters using an IFFT.

rithm eliminates the need for twiddle factors, which can provide greater efficiency for a large number of channels.

6. REFERENCES

REFERENCES

- [1] fredric j. harris, *Multirate Signal Processing for Communication Systems*. Upper Saddle River, NJ: Prentice Hall PTR, 2004.
- [2] P. S. Crescini, "Evaluation of channelizer architectures for space applications," Master's thesis, San Diego State University, 2014.
- [3] f. harris, E. Venosa, X. Chen, and M. Renfors, "Cascade linear phase recursive half-band filters implement the most efficient digital down converter," *Analog Integrated Circuits and Signal Processing*, vol. 73, no. 2, pp. 531–543, 2012. [Online]. Available: <http://dx.doi.org/10.1007/s10470-012-9905-9>
- [4] fredric j. harris, "A guided tour and comparison of fft algorithms: Cooley-tukey and good-thomas," July 2013.