An efficient channelizer tree for portable software defined radios

Fred Harris, Elettra Venosa, Xiaofei Chen & Chris Dick

annals of telecommunications - annales des télécommunications

ISSN 0003-4347

Ann. Telecommun.
DOI 10.1007/s12243-013-0354-y

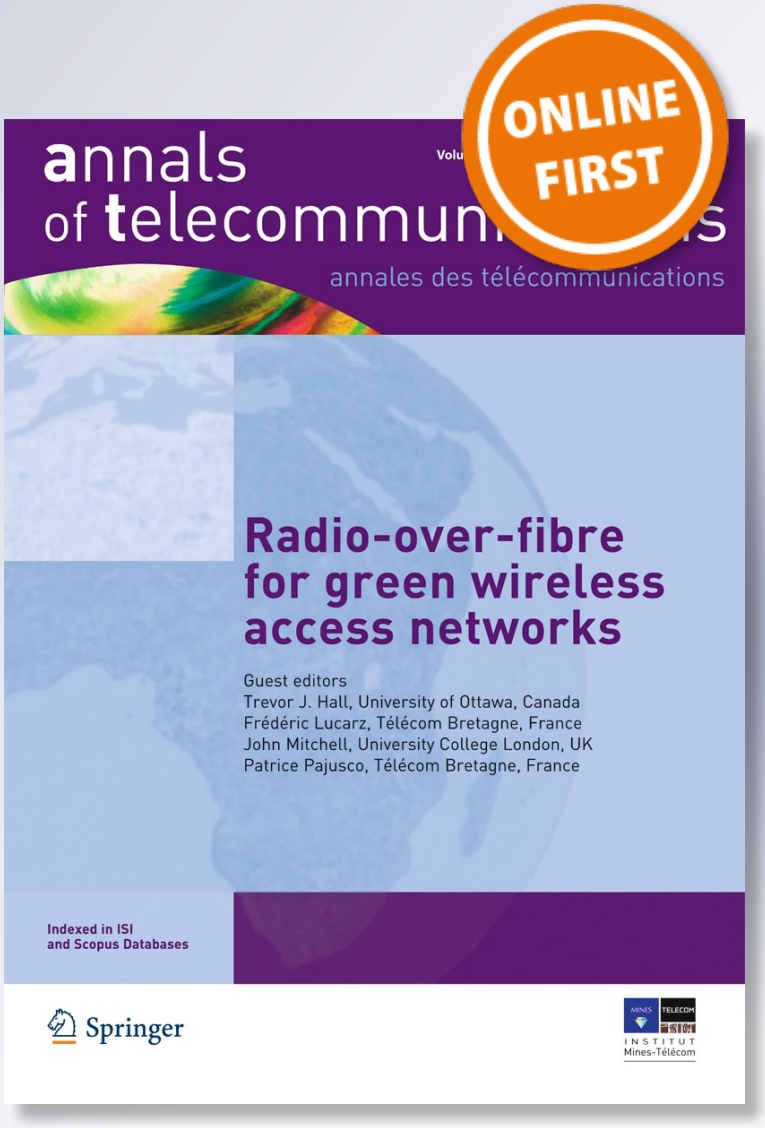

Springer

Ann. Telecommun.
DOI 10.1007/s12243-013-0354-y

An efficient channelizer tree for portable software defined radios

Fred Harris · Elettra Venosa · Xiaofei Chen · Chris Dick

Received: 14 March 2012 / Accepted: 9 January 2013

Abstract Power consumption is one of the most critical issues in the portable software-defined radio devices. A software radio receiver has the need to downconvert, bandwidth limit, and downsample a single narrowband channel from a span of frequencies in the Nyquist zone collected by the input analog to digital converter. In this paper, we present two techniques that perform the receiver function more efficiently than the standard Gray chip architecture formed by its direct digital synthesizer (DDS), and two stages of downsampling with a cascade integrator comb (CIC) filter and a pair of half-band filters. We compare the workload of this conventional architecture to two new architectures by applying them to the task of extracting a single, 30 kHz wide, channel from a 30 MHz band sampled at 90 MHz. One proposed structure replaces the CIC filter with a 10-stage cascade of 2-to-1 downsampling half-band filters with successively narrower transition bandwidths. In the second proposed structure, the DDS is moved to the output of the filtering stages which perform a sequence of 2-to-1 downsampling operations in half-band filters that perform incidental spectral translation by aliasing. We enlarge the set of half-band filter that reside at 0 and fs/2, to also include the Hilbert transform half-band filters residing at ±fs/4. At every stage in the cascade, the selected band resides in one of the four half-band filters. The 2-to-1 downsampling with that filter reduces the bandwidth and aliases the desired center through a sequence of known center frequencies. The desired channel is recovered from the output of the final stage by a complex heterodyne applied (at the low output rate) to obtain the desired spectral shift to base band. The paper provides a detailed workload analysis of the proposed structure along with simulation results that prove its effectiveness.

Keywords Polyphase filter bank · Down converter channelizer · Software defined radio

1 Introduction

There are many ways for selecting and implementing a process to downconvert and extract a single narrowband channel from a set of narrow bandwidth channels distributed over an observed frequency span. The need to minimize the power consumption in the portable software radios pushes us to work in the direction of finding the most efficient solution for accomplishing this task. An accurate analysis of the many possible options is required to understand which one provides the minimum power consumption.

In this paper, the authors present two flexible options based upon a cascade of half-band filters that select and downconvert a single channel from a large span of bandwidth with very low computational workload. The simulations and the workload analysis are based on a practical scenario of channels having bandwidth of 30 kHz residing in the frequency band between 0 and 30 MHz that has been sampled at a 90 MHz rate.

The standard processing structure to downconvert, filter, and downsample a selected channel is shown in Fig. 1. Here, we see a complex heterodyne, followed by a pair of low-pass filters and a downsampler that reduces the

F. Harris (✉) · E. Venosa · X. Chen
San Diego State University, San Diego, CA 92182-1309, USA
e-mail: fred.harris@sdsu.edu

E. Venosa
e-mail: evenosa@projects.sdsu.edu

X. Chen
e-mail: chenxiaofei_sdsu@yahoo.com

C. Dick
Xilinx Corp, San Jose, CA 95124, USA
e-mail: chris.dick@xilinx.com

Published online: 08 February 2013

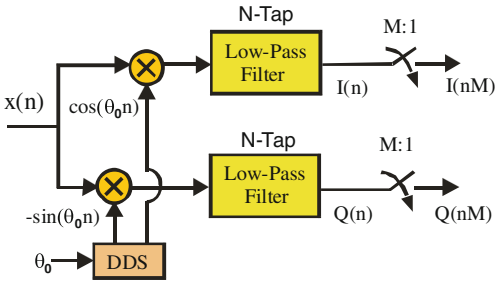

Fig. 1 Conventional digital down converter

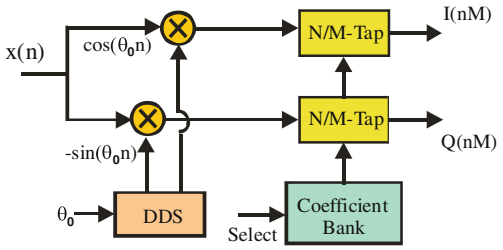

Fig. 3 "Gatling Gun" implementation of M-to-1 down-sampling low-pass FIR filters in digital down converter

sampling rate commensurate with the reduced signal bandwidth. The output data from the quadrature digital down converter (DDC) are complex so the low-pass filter, having real impulse response, is implemented as two identical filters, each processing one of the quadrature time series [3]. This architecture is not efficient when applied to input signals with a large ratio of sample rate to bandwidth. In this case, its computational workload is high due to the large number of coefficients required to implement the filter as well as having to apply the heterodyne to the signal at the high input sample rate.

The standard response to this inefficiency is to interchange the tasks of filtering and resampling so that only the samples destined to be preserved by the resampling process are computed. This makes great sense since it would be foolish to compute output samples and then discard them. The polyphase partition of the low-pass filters following the quadrature downconversion is shown in Fig. 2. Since the polyphase arms are accessed sequentially by their input commutators, the polyphase partition can be implemented as shown in Fig. 3 by a single N/M tap filter per path that accesses the successive weight sets. We liken this version to a Gatling gun model where, as in the gun, successive weights are delivered to the processing engine as successive input samples are delivered to the engine from the quadrature mixers. The N/M-Tap filter in this version of the process is implemented in the partial sum accumulator form. In this dual form of the tapped delay filter, the distributed accumulators store products of the input samples rather than the input samples.

Another technique to reduce the computational workload for a DDC is to perform the bandwidth reduction and sample rate change in a two-stage filter. Figure 4 is an example of this option. Here too, the reduction in workload is impressive because the first stage's effort is significantly reduced by virtue of the wider transition bandwidth required of the filter design and the second stage effort is similarly reduced by operating at a lower input rate. Both filters in the cascade can be implemented in the polyphase form shown in Fig. 2.

One well-known version of this cascade option, shown in Fig. 5, uses the cascade integrator comb (CIC) filter for the first stage to effect a large bandwidth reduction and sample rate reduction and the second process is a pair of half-band filters that corrects the in-band spectral droop of the CIC and performs a final bandwidth and sample rate reduction at the reduced sample rate from the resampling CIC filter. The initial attraction of this version of the cascade is that the CIC is implemented without multipliers.

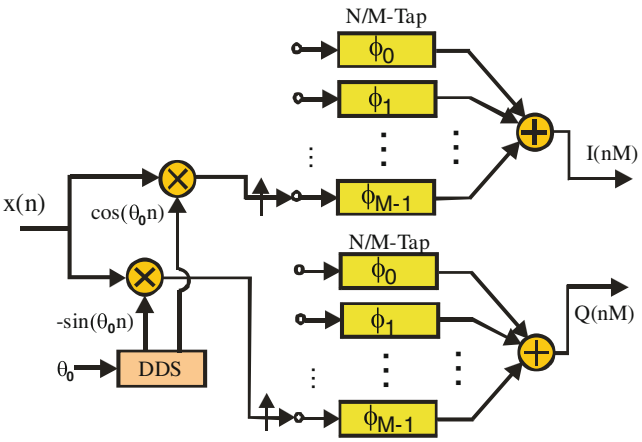

Fig. 2 Polyphase partition of M-to-1 down-sampling low-pass FIR filters in digital down converter

2 Polyphase down converter

We are well aware that when we apply the Noble identity to an M-path filter, we interchange the operations of filtering and resampling. By this exchange, we embed the resampler in the filter and significantly reduce the work load by

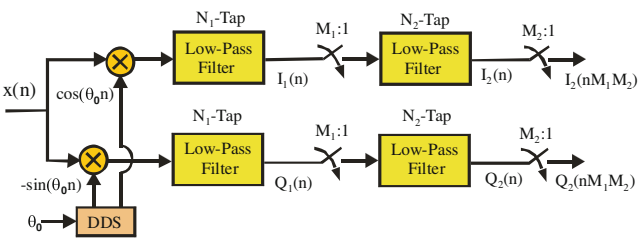

Fig. 4 DDC with two-stage filter and resampler

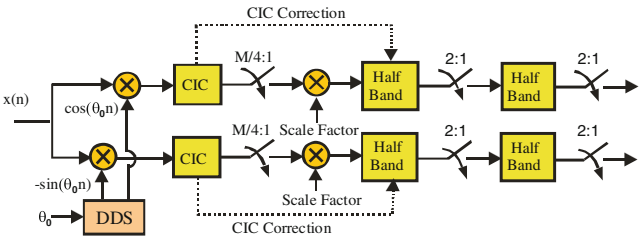

Fig. 5 DDC with CIC first-stage filter and resampler

operating the filter at the output rate rather than at the input rate. We can further reduce the workload of the DDC by applying another filtering relationship to the DDC known as the equivalency theorem. This relationship tells us that an input downconversion followed by a baseband low-pass filter is equivalent to a band-pass filter followed by an output downconversion. Pursuing the goal of reducing the total workload of the DDC system, we can interchange the operations of downconversion and filtering and shift the complex heterodyne to the output of the filter. The reduction in workload occurs when we move the complex heterodyne to the output of the downsampler and only downconvert the reduced number of output samples at the output rate. This sequence of transformations is shown in Fig. 6.

We can now reorder the band-pass filter and the resampler to convert the filters to their polyphase equivalents, shown in Fig. 7, which also operate at the reduced output sample rate. Our final trick is to select the center frequency of the band-pass filter to be a multiple of the output sample rate. For this condition, the downsampling aliases the center

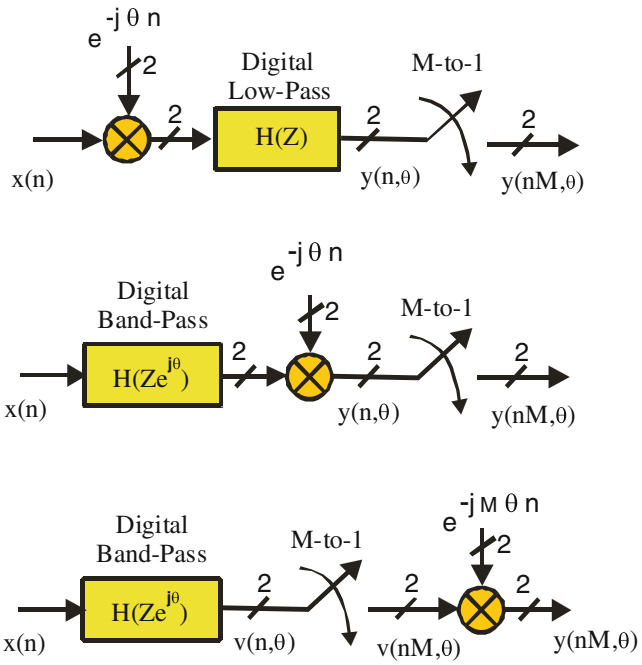

Fig. 6 Sequence of transformations move down converter from input port at input rate to output port at output rate

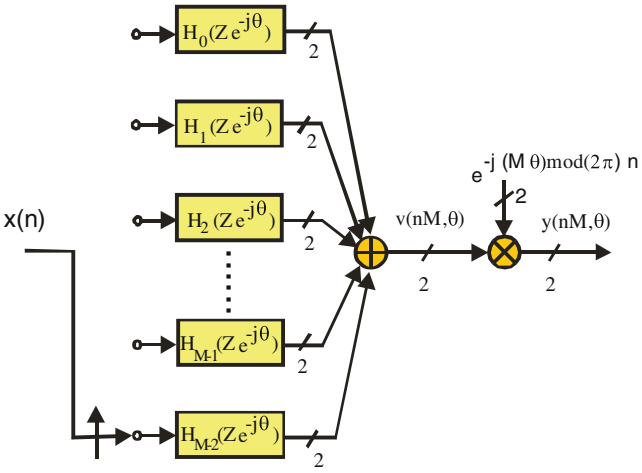

Fig. 7 Polyphase band-pass filter embedded in DDC

frequency to DC and the related heterodyne defaults to a unity-valued scalar which is removed from the signal processing path. In addition, the rotators in each path become the same value and can be factored out of the filter path and applied as a single complex rotation at the output of each path. This option is shown in Fig. 8 [1]. The unique property of this configuration is that the real input samples delivered to the polyphase filter do not become complex till they pass through the complex rotators following the filter in each path. This means that we only require a single filter, not a pair of filters as required by the earlier implementations of the DDC. This is actually a remarkably efficient implementation but we are seeking an even more efficient option.

3 Proposed structure 1

We propose two enhanced versions of the DDC as variations of the DDCs discussed in the introduction. In the first

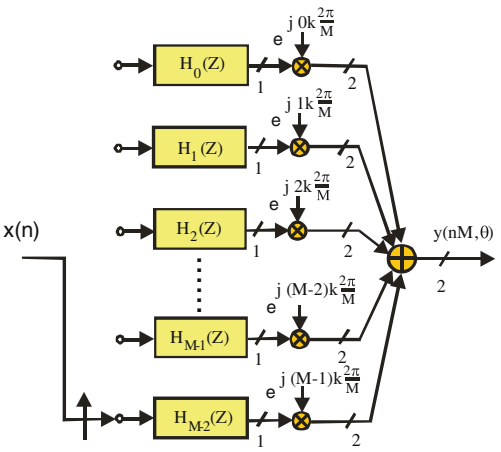

Fig. 8 DDC polyphase band-pass filter with center frequency at a multiple of output sample rate

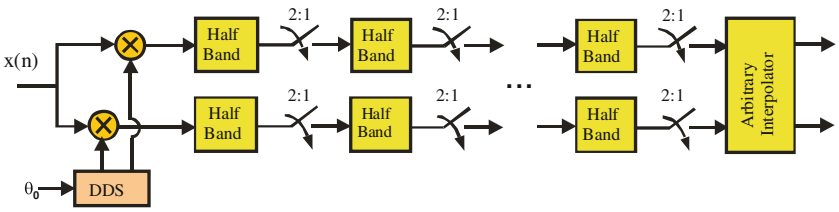

Fig. 9 Cascade of K 2-to-1 down sample half-band filters

option, we extend the design illustrated in Fig. 4 from a cascade of two filters to a cascade of K half-band 2-to-1 downsampling half-band filters. This option is shown in Fig. 9.

This cascade containing 10 half-band filters performs a 1,024 to 1 bandwidth reduction to 30 kHz with a sample rate reduction from 90 MHz to 87.8906 kHz. The final processing block, an arbitrary interpolator, will change this output sample rate to a convenient output rate for subsequent processing. For instance, we may want to have an output sample rate of two samples per symbol (60 kHz) or some other sample rate related to legacy postprocessing algorithm requirements.

Exact half-band filters with an odd number of taps are very efficient engines since, except for the center weight, the even indexed coefficients are zero and the remaining weights are even symmetric permitting a folded version of the filter to share the common weight on either side of the center sample. The filters are designed to suppress, by 100 dB, the 30 kHz wide band centered at the half sample rate which aliases onto the desired DC centered 30 kHz band by the 2-to-1 downsampling operation.

Because the 30 kHz band is such a small fraction of the 90 MHz sample rate, the first three filters are implemented with a pair or repeated zeros at the half sample rate. Their impulse response weights are [1, 2, 1]/4 which requires only two add and two shift operations to implement. The fourth filter in the chain also requires three weights but the coefficients were adjusted to widen the stop-band bandwidth by separating the zeros near the half-sample rate. The six remaining filters in the chain were designed as true half-band filters by the Remez algorithm and Vaidyanathan's half-band filter trick [2]. Figure 10 presents the spectra of the five filters along with their stop band masks to illustrate the performance of their stop band characteristics.

The impulse response and frequency response of the tenth filter in this chain is shown in Fig. 11. This filter has 17 coefficients with eight of the even indexed coefficients equal to zero and the center coefficient equal to 0.5. There are eight odd indexed coefficients with even symmetry for which if implemented as a folded filter would require four multiplies per output sample or two multiplies per input sample when operating at its half sample rate. As mentioned, earlier, the first three filters in the chain have trivial weights of [0.25, 0.5 0.25] which place repeated zeros at fs/2. These filters can be implemented with two adds and shifts. Table 1 lists the filter lengths and the work load required to implement them. Table 2 lists the filter numbers

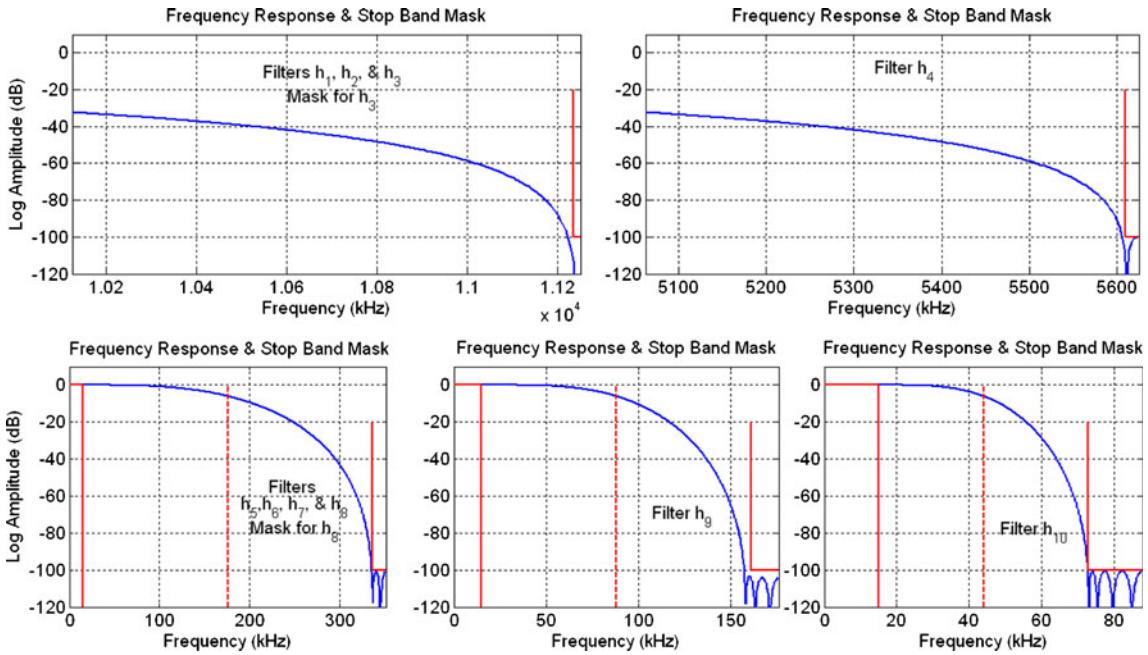

Fig. 10 Frequency response and stop band mask for filters h_3, h_4, h_8, h_9, and h_{10}

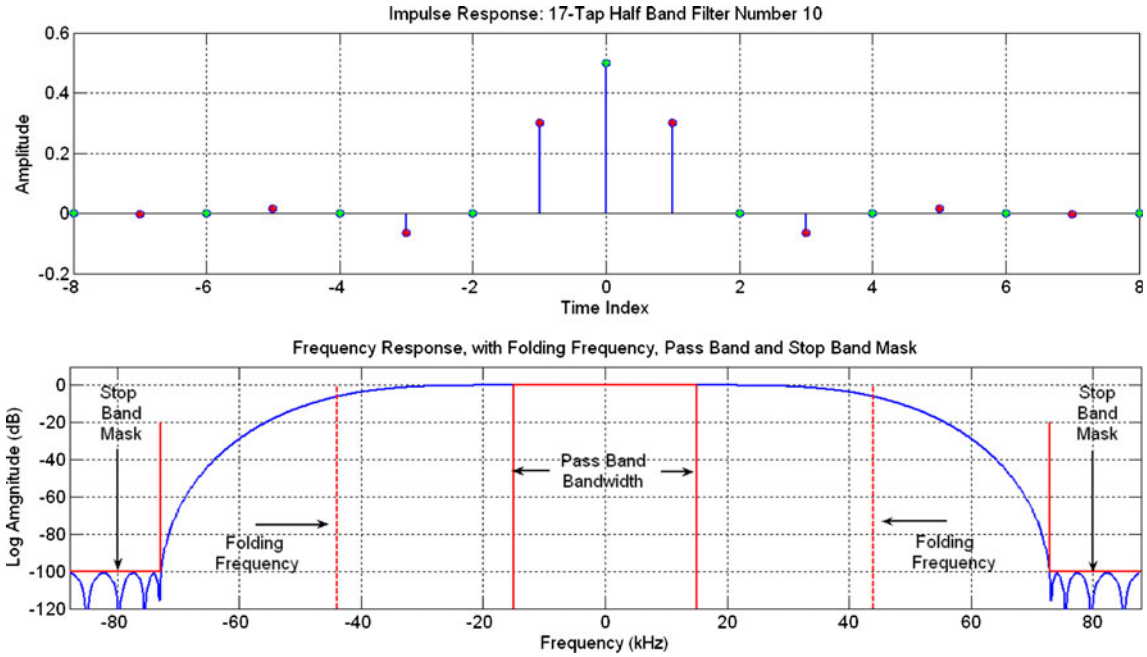

Fig. 11 Impulse response and frequency response of 2-to-1 downsampling half-band filter h_{10}, in cascade

and the number of multiplies and adds per stage and the number of multiplies and adds referenced to the input. For instance, filter 2 has the same workload as filter 1 but operates at one half of the input rate and filter 3 also has the same workload as filter 1 but operates at one fourth of the input rate. The total workload for the 10 stages of filtering is seen to be 2.13 adds and 0.19 multiplies per input sample. Of course there are two of these filters in the DDC and we still have to account for the two multiplies per input sample in the quadrature down converter. Doing so, we learn that the total workload for this version of the DDC is less than 2.4 multiplies and 4.3 adds per input sample point.

It is appropriate at this point to compare the computational workload of our proposed half-band filter cascade with the ubiquitous DDC formed by the CIC and half-band filter pair shown earlier in Fig. 5. Figure 12 presents the frequency response of a four-, five-, and six-stage CIC filter along with 100 dB stop band spectral masks corresponding to the output bandwidth equal to one quarter of the output sample rate.

From this figure, we see that a six-stage CIC is required to obtain the 100 dB suppression of the band that aliases to baseband with the 2-to-1 downsample. To obtain a 256-to-1 downsample from 90 MHz to 351.56 kHz with a 30 kHz output bandwidth and with the same 100 dB dynamic range of our half-band filter chain, we require six CIC stages. The gain of each stage is 256 which after six such stages becomes $(256)^6$ or 2.8×10^{14} which would require 48 bits of growth in the CIC integrators. If we assume 16-bit input data, the bit width of the six integrators would have to be 16+48 or 64.

In total, the six integrators in both the I and Q paths would be circulating 768 bits per input sample which if converted to the 20-bit width (assuming 5 dB/bit) required of the arithmetic in the half-band filters proves to be same number of bits manipulated in 38 arithmetic operations per input sample. The number of operations for our first proposed I–Q half-band filter chain is on the order of 4.3 adds and 0.38 multiply. If we count a multiply as being equivalent to 20 adds our cascade requires approximately twelve 20-bit adds per input sample which represents a workload 1/3 the CIC chain. This suggests that a 3-to-1 power savings can be had by replacing even the very efficient cascade CIC filter chain with the cascade half-band filters.

This arithmetic savings does not include the workload of the pair of half band filters in each path following the CIC filters required to correct their in-band spectral droop.

Table 1 Linear phase half-band FIR filters in DDC cascade with number of taps and number of arithmetic operations at output sample rate

Filter numbers	Number of taps	Operations
h1-to-h3	5	2-A, 0-M, 1-Shft
h4	5	2-A, 1-M,0-Shft
h5-to-h8	9	4-A, 2-M, 0-Shft
h9	13	6-A, 3-M, 0-Shft
h10	17	8-A, 4-M, 0-Shft

Table 2 Filter number in FIR cascade, number of taps in filter, and number of arithmetic operations referenced to input

Filter number	1	2	3	4	5	6	7	8	9	10	Total
Number of taps	5	5	5	5	9	9	9	9	13	17	–
Adds	2	2	2	2	4	4	4	4	6	8	–
Multiplies	0	0	0	1	2	2	2	2	3	4	
Adds referenced to input	2/2	2/4	2/8	2/16	4/32	4/64	4/128	4/256	6/512	8/1024	2.13
Multiplies referenced to input	0/2	0/4	0/8	1/16	2/32	2/64	2/128	2/256	3/512	4/1024	0.19

4 Modified polyphase down converter

In our introduction, we presented the M-path polyphase downconverter and presented its structure in Fig. 7. We return to this structure, examine it, and suggest an interesting version of it, which will become the building block of our second proposed DDC. In its most common incarnation, an M-path polyphase channelizer can downconvert and downsample M equally spaced, fixed bandwidth channels. We saw in Fig. 7 that its structure is formed by an M-port commutator, an M-path partitioned low-pass prototype filter, and a set of M complex rotators.

In this engine, the commutator delivers M consecutive samples to the M input ports of the M-path filter performing the signal sample rate reduction which causes M spectral folds in the frequency domain. With an output sample rate of f_S/M, all M multiples of the output sample rate alias to base band (DC). The alias terms in each arm of the M-path filter exhibit unique phase profiles due to their distinct center frequencies and the time offsets of the different downsampled time series delivered to each commutator port. The partitioned M-path filter aligns the time origin of the sampled data sequences delivered by the input commutator to a single common output time origin. This task is accomplished by the all-pass characteristics of the M-path partitioned filter that applies the required differential time delay to the individual input time series. Finally, the phase rotators perform the equivalent of a beam-forming operation; the coherent summation of the time aligned signals at each output port with the selected phase profile [4].

The phase-coherent summation at the output of the M-path filters extracts a selected alias residing in each path by constructively summing the selected aliased frequency components located in each path, while simultaneously destructively canceling the remaining aliased spectral components. The phase rotators extract, in each arm, from the myriad of aliased signals only the alias with the particular matching phase profile.

We note that the four-path partitioned filter case has a unique characteristic: the four-phase rotators, to be applied to extract any of the four channels, are trivial combinations of ± 1 and $\pm j$. This is a consequence of the fact that the center frequencies of the filters' aliased copies are centered on multiples of the quarter sample rate. In this case, the phase rotators applied in each arm of the channelizer for extracting the selected channel to base band default to unitary values, hence do not contribute to the workload of the processing chain.

In normal channelizer operation mode, the bandwidth of the prototype filter to be partitioned into the four paths, would be one quarter of the input sample rate. In our modified channelizer, we will use a true half-band FIR filter as our low-pass prototype filter. Doubling the prototype bandwidth requires us to double the output sample rate. We perform a four-path partition of the half-band filter and modify its structure so that it performs 2-to-1 downsampling. This is an interesting extension of our two-path half-band filter. The modified form of the four-path filter is shown in Fig. 13. Note that the input commutator delivers input samples to two paths simultaneously. Because of the

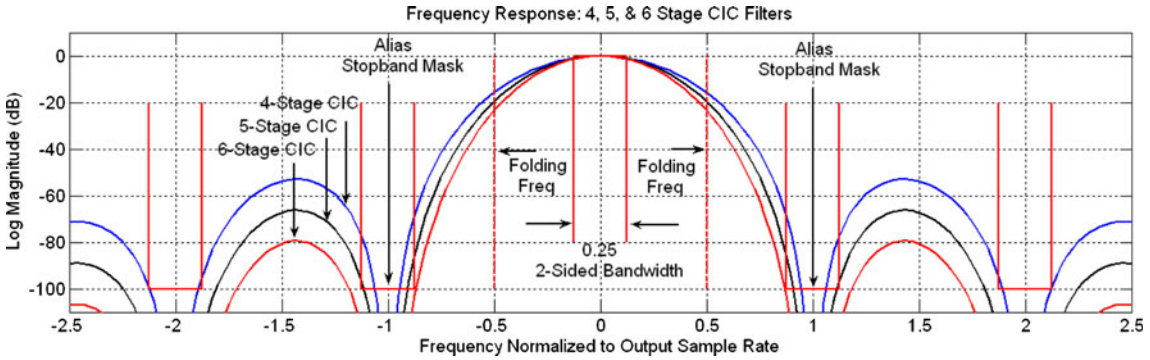

Fig. 12 Frequency response of four-, five-, and six-stage CIC filters showing stop band spectral mask locations

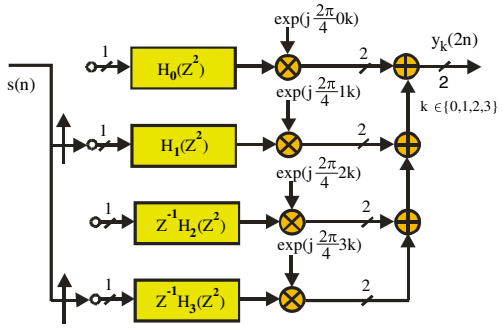

Fig. 13 Four-path, 2-to-1 downconverter channelizer

delay in the lower two paths, the upper two paths use the current input samples to compute the current output while the lower two paths use the current inputs to compute the next output. The four versions of the half-band filter available from this structure reside at multiples of fs/4. The two spectra that reside at DC and fs/2 are the pair to which we have access in a normal quadrature mirror filter. The two added center frequencies at ±fs/4 are seen to be Hilbert transform filters. As we will see in a moment, path 2 is removed from the block diagram.

When a half-band filter is downsampled 2-to-1, the transition band edges alias into its pass band. We design the half-band filters with transition bandwidth symmetrically positioned about the quarter sample rate extending from 0.125 to 0.375 so that the 2-to-1 downsampled bands between ±0.125 about the four possible band centers are alias free. Thus, each of the four center frequencies offers an alias free band representing a quarter of the available bandwidth. The alias free band is widened slightly to accommodate a selected signal bandwidth located at boundaries of ±0.125. The impulse response and frequency response of a 23 tap half-band filter is shown in Fig. 14. Seen here is the folding of the transition band due to the 2-to-1 downsampling as well as the four spectral locations available from the four-path filter. Table 3 shows the partition of the half-band filter weights into the four paths of the four-path filter. Note that path 2 weights are zero and this path can be omitted from the block diagram. Table 4 lists the four sets of trivial phase rotators that combine the outputs of the four path filters to obtain the filter output at the center frequencies 0, fs/4, 2fs/4, and 3fs/4. Also note the weights in path 1 and path 3 are mirror images so that this filter can also be folded to reduce the multiply rate by a factor of 2.

Table 3 Half-band filter coefficients in four-path filter

0	0	0	0.5	0	0
−0.0004	−0.0107	−0.0828	0.3097	0.0314	0.0028
0	0	0	0	0	0
0.0028	0.0314	0.3097	−0.0828	−0.0107	−0.0004

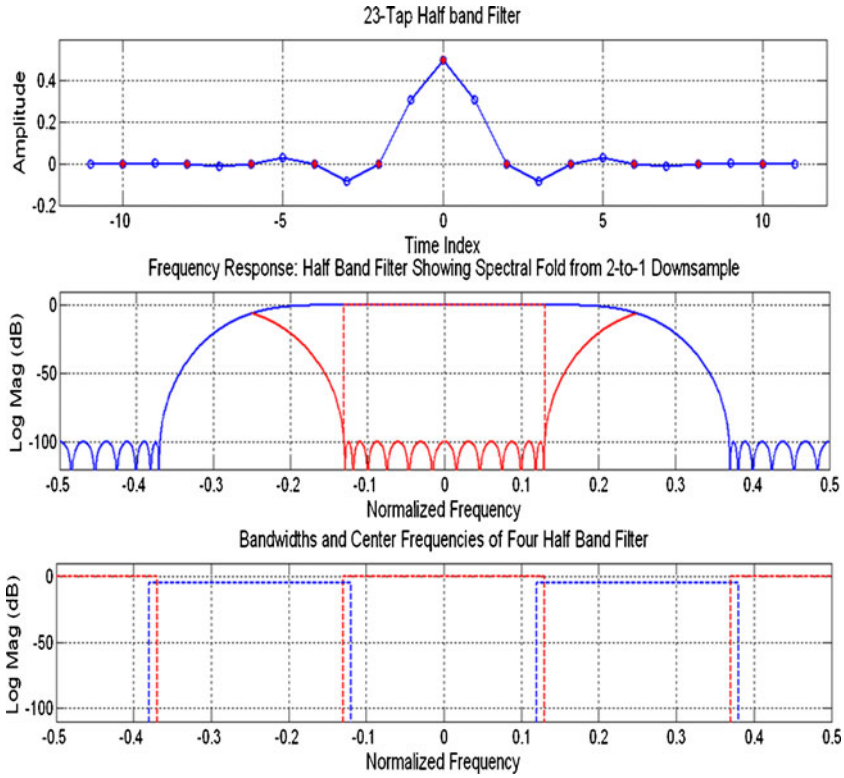

Fig. 14 Impulse and frequency response with aliased transition band due to 2-to-1 downsample and four locations from four-path channelizer

Table 4 Phase rotators combining output of four paths

	0	fs/4	2fs/4	3fs/4
Path 0	+1	+1	+1	+1
Path 1	+1	j	−1	$-j$
Path 2	–	–	–	–
Path 3	1	$-j$	−1	$+j$

5 Proposed structure 2

We now propose the second of our enhanced version of the DDC as variations of the DDCs discussed in Section 1. The proposed structure, depicted in Fig. 15, is a cascade of polyphase stages each one composed of a modified four-path channelizer performing 4-to-2 downsampling. The number of stages we use here is 10 as it was for the first proposed cascade of half-band filters.

This option is an unusual DDC because the selected band to be downconverted is not downconverted on the way into the cascade filter. Rather the known center frequency of the selected band is intentionally aliased to a succession of different center frequencies as a result of the successive 2-to-1 downsampling operations. Since the initial center frequency is known the sequence of new centers is also known. The known center frequencies are located in one of the four filter bands and at each stage, we select the filter in which the center is located to perform the bandwidth reduction and sample rate change. We know the final center frequency to which the selected channel center has aliased as the result of the multiple 2-to-1 downsampling operations. We finish the downconversion task by applying, at the output sample rate, a quadrature heterodyne for the final translation to baseband. In this option, the spectral translation is applied at the low output rate rather than at the high input rate.

Note that due to the effect of the sample rate halving at each stage, the relative position of the selected signal center frequency is doubled when normalized to the new output sample rate. As an example, we track the first four successive locations of a signal located at normalized input frequency of 0.1 which places it in the DC bin, or Bin-0 (−0.125 to +0.125) of the four-path filter. After the first half-band filter and 2-to-1 downsample, it has aliased to the normalized frequency 0.2 which places it in Bin-1 (+0.125 to +0.375). Following the second half-band filter and 2-to-1 downsample, it has aliased to 0.4 which is in Bin-2 (+0.375 to +0.625). The third half-band filter and 2-to-1 downsample places it at 0.8 which is in Bin-3 (+0.6255 to +0.875). The fourth half-band filter and 2-to-1 downsample places it at 1.6 or 0.6 modulo(1) which is back in Bin-2 (+0.375 to +0.625).

Note that, we do not have to shift the selected channel to base band at each half-band iteration but do so at the end of the alias chain. In each stage, we only select the channelizer channel that contains the aliased center frequency.

The true half-band filter is used in the four-path channelizers because it has the particularly attractive property that, except for the sample value at the origin, all the even indexed coefficients in its impulse response have zero value. The reduced number of non-zero filter taps implies a commensurate decrease in the workload for each stage. In our structure, we design the half-band filters in a way such that, as a consequence of the four-path partitioning, the pass band of the filters' aliased copies only overlap by 30 kHz, which is exactly the gap between the center frequencies of two input signals. That overlapping bandwidth is sufficient to guarantee that no channel will be lost during the processing.

Because the pass band center frequency of the last stage half-band filter is not likely to have been aliased to zero frequency, a complex heterodyne followed by a final low-pass filter is included in this design.

6 Simulation results

In this section, we present the simulation results regarding the 10-stage channelizer tree derived in the previous sections. The signal selected to be downconverted is the 600th center frequency of the 1,000 frequencies spanning the 30 MHz signal bandwidth in 30 kHz increments starting from 0.

Figure 16 presents the spectrum of the real input signal. Since the spectrum is Hermitian symmetric, only the positive frequencies of the spectrum are plotted. Because it is composed of 1,000 signals and it is difficult to distinguish each of them, the lower subplot gives an enlarged view of the upper subplot between 17.9 and 18.1 MHz. In this subplot, seven of the thousand 30 kHz spaced signal bands composing the input spectrum are visible.

Figure 17 shows the output spectra of the last four stages of this second proposed structure. The sampling frequency of the other stages is too high to distinguish the output spectra. A dotted line is used to indicate the center

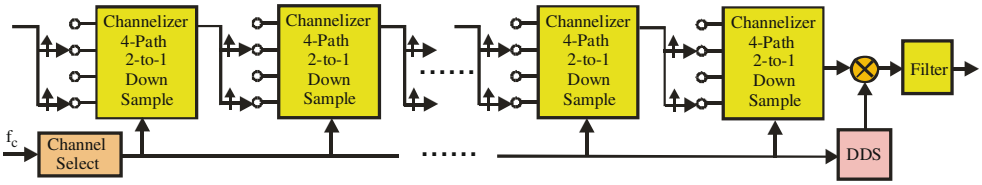

Fig. 15 Second proposed 10 stages half-band filter structure

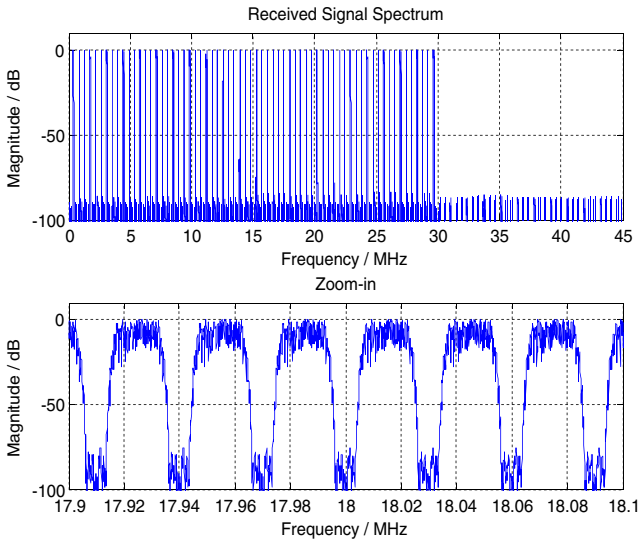

Fig. 16 Spectrum of the input signal

frequency of the signal we have selected. In the last subplot of Fig. 17, we see three of the input spectra residing in base band. Only one of them is the desired one. The complex heterodyne and the low-pass filter will shift it to baseband and select it from the three available spectra. In Fig. 18, the output signal spectrum from the complete processing chain is shown. In particular, the upper subplot of this figure is coincident with the last subplot of Fig. 18. It represents the output of the last channelizing stage. The middle subplot of Fig. 18 represents the heterodyned spectrum with superimposed frequency response of the following low-pass filter. In the lower subplot of this same figure, we have the spectrum of the final output of this structure.

7 Workload analysis

In this section, we calculate the total workload of the proposed processing chain. Let $P(k)$ denote the number of non-zero coefficients in the kth half-band filter. This number is 16 for the first six stages and this number increases as the single channel overlap of the half-band filters becomes a more significant fraction of the successive output sample rates. The final processing block of the chain includes a complex heterodyne and a pair of final bandwidth selection filters of length L. Each successive half-band filter operates at half the rate of the previous stage. The first stage chain processes real samples and in the most general case, the remaining stages process complex samples.

Remember that with half the coefficients delivering two inputs to compute one output, the number of multiplies per input to the filter is $P(1)/4$. We will assume the final stage in the chain that isolates the selected channel is a filter with L weights. If we choose to fold the coefficients about their midpoint to take advantage of their symmetry, the number of multiplies per input can be further reduced to P(1)/8. For this scenario, the total number of multiplies reflected to the input sample rate of the structure without coefficient folding is seen to be:

$$W = \frac{P(1)}{4} + 2\left(\frac{1}{2}\frac{P(2)}{4} + \frac{1}{2^2}\frac{P(3)}{4} + \ldots + \frac{1}{2^9}\frac{P(10)}{4} + \frac{1}{2^{10}}L\right) \quad (1a)$$

With coefficient folding, the number of multiplies per input is:

$$W = \frac{P(1)}{8} + 2\left(\frac{1}{2}\frac{P(2)}{8} + \frac{1}{2^2}\frac{P(3)}{8} + \ldots + \frac{1}{2^9}\frac{P(10)}{8} + \frac{1}{2^{10}}\frac{L}{2}\right) \quad (1b)$$

In the nonfolded version of the filter, the number of adds is equal to the number of multiplies and in the folded form, the number of adds is twice the number of multiplies.

The filter lengths required for 10 stages of half-band filtering with successively reduced transition bands was found to be the values [19, 19, 21, 21, 21, 21, 23, 25, 31, 65]. The non-zero coefficients at each stage is [10, 10, 10, 10, 10, 12, 12, 14, 16, 32] which when substituted in Eq. (1a, 1b) present the result of approximately 7.70 (or 3.85 if we fold the coefficients).

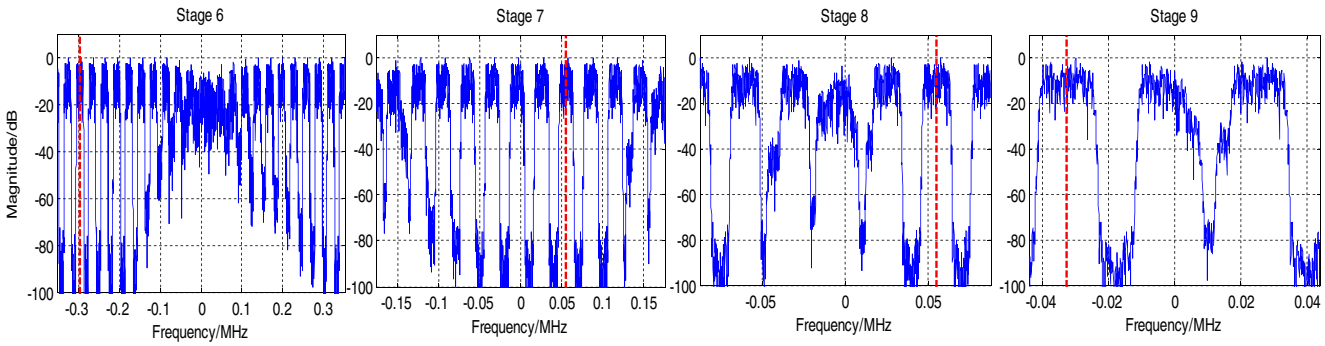

Fig. 17 Outputs at the last four stages of the processing chain

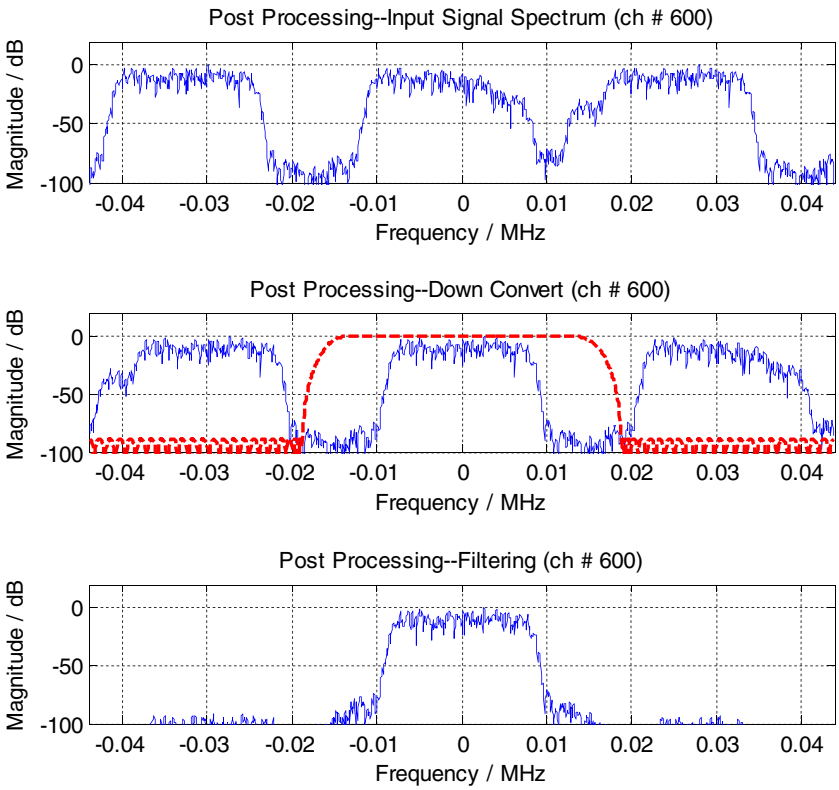

Fig. 18 Output signal of the proposed structure after the complex heterodyne and the final filtering process

Had we assumed all the filters were of equal length, we would have an upper bound shown in Eq. (2) which when evaluated with P=10 would offer an upper bound of 7.75 (or 3.88 with folding, a value within 1% of the actual value.

$$\frac{\text{ops}}{\text{input}} = \frac{P}{4}\left[1+1+\frac{1}{2}+\frac{1}{4}+\cdots\right]+2\frac{L}{2^{10}} < \frac{3P}{4}+\frac{120}{512} \simeq 0.75\times P+0.25 \quad (2)$$

Table 5 compares the relative workload for the three filters we have been comparing, the CIC, the first proposed option, the second proposed option and for thoroughness, an M-path polyphase filter. We note that the CIC has, by far, the largest number of effective additions per input sample. Remember that we have converted the multiple, very wide additions of the CIC input accumulators, into an equivalent number of 20-bit additions.

Table 5 Effective number of multiplies and adds for different implementations of digital downconverter with 1,000-to-1 downsampling

	Mixer multiplies	Filter multiplies	Filter adds
CIC	2	0	38
Two-path HB	2	0.38	4.3
Four-path HB	0	3.88	7.76
M-Path	2	8	8

We will leave it to reader to interpret how many 20-bit additions are equivalent to each multiply. An ASIC designer will compare real estate for the two operations and a programmer will not care since they already have the multiplier and not using it will be a waste of resources. Most likely, an energy dissipation budget is the correct way to compare the four options. We will also leave open the question if there are other good mixes of cascade filters that would offer improved workload or energy dissipation numbers.

8 Performance comparison

Till now, we have emphasized energy budget differences between the CIC filter and the cascade half-band filter chain. We now compare the ability of the two filter options to perform their tasks of bandwidth and sample rate reduction with minimal reduction in signal fidelity. Every DSP filtering operation contributes signal distortion associated with spectral gain and phase deviation as well as processing noise related to finite bit width arithmetic. The standard way to measure the cumulative contribution of distortion and noise is to probe the filtering process with a communication signal designed to have, when properly demodulated, zero intersymbol interference.

Constellation points in two-space initiate the modulation process and matching constellation points are extracted from the demodulation process. The perfect modulation and

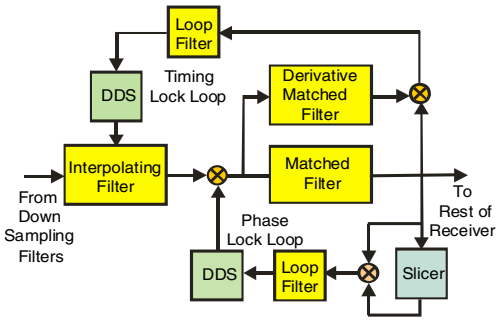

Fig. 19 Demodulator process following both digital downsampling filter process

demodulation process would result in single-value constellation points. The distortion and computational noise results in small random displacements of constellation points centered about the nominal constellation point positions. We use the second central moment, the variance of the constellation cloud, as a convenient measure of the distortion and noise contributed by the filtering process. This measure is called the error vector magnitude (EVM): small EVM is good and large EVM is undesirable.

The proper demodulation process entails the phase alignment of the constellation points about their nominal position and the correct time positioning of the sampled data point of the complex wave shape envelope. Figure 19 shows the block diagram of the receiver structure that phase aligns and time aligns the sample position of the detected constellation point. This process uses a phase lock loop to rotate the selected sample points and uses an interpolator to time shift the correct sample value of the complex envelope to one of the clock sample points. The interpolator serves a second function: that of changing the sample rate obtained from the digital downsample circuitry, 90 MHz/1024, or 87.8906 kHz, to a multiple of the modulation symbol rate of 15 kHz. We selected 90 kHz, six times the output symbol rate, as a convenient output sample rate. We applied this demodulation process to the output of the 1024-to-1 downsampled half-band data signal and to the 256-to-1 CIC filter followed by two 2-to-1 downsampled half-band filter and a compensation filter that removes the spectral curvature of the CIC filter's main lobe response.

The left subplots of Fig. 20 show the constellations obtained from the demodulation processes applied to the two-filter options. Every sixth sample is identified by a red sample point: these points are the constellation points of the demodulation process. The right subplots are zooms to the immediate vicinity of the constellation cluster centered at the order pair $1+j1$. We see that the radius of the half-band cluster is approximately one part in a 1,000 while the radius of the CIC cluster is approximately 2.5 parts in 1,000. The EVM values were computed to be −32.44 and −30.01 dB, respectively; a difference of 2.4 dB! What a surprise! Why is there a difference? Both filters are performing the same task using different techniques! Ah, it is the difference in the

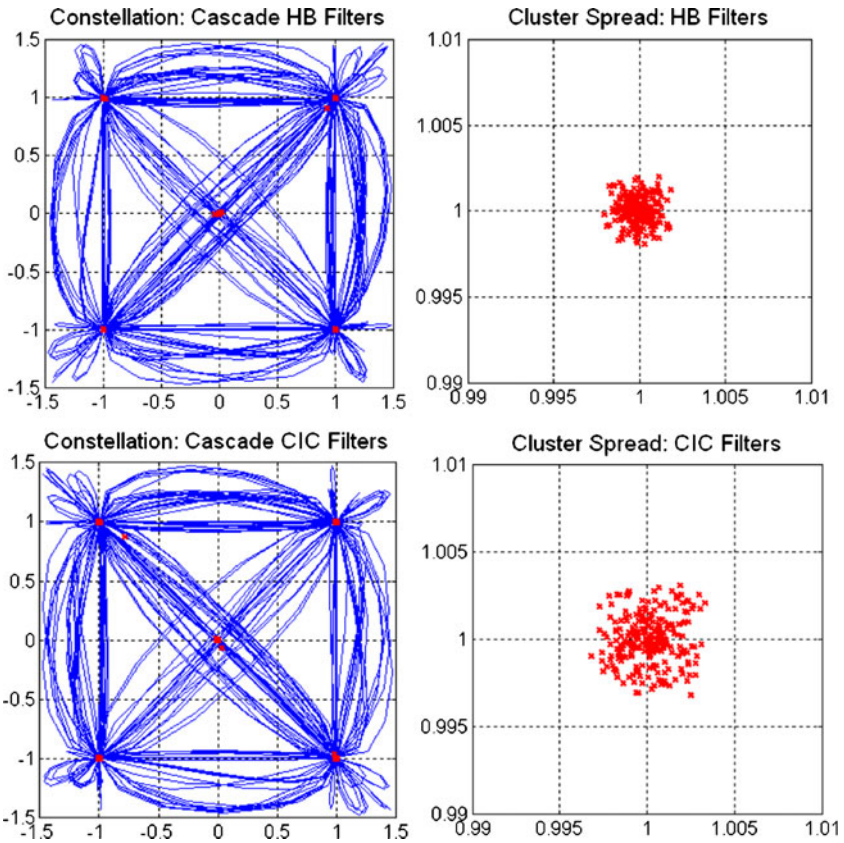

Fig. 20 QPSK constellations and detailed constellation clusters for downsampling half-band filters and CIC filters

techniques that is responsible for the difference in performance. We observe that both filter options processed the same input data stream and the same receiver structure post processed the outputs of the two-filter options. Since the only difference in the two processing paths was the downsampling filter options, we must attribute the performance difference to the downsampling filter processes.

This is important! When the half-band filter performs a 2-to-1 downsample a single cleared, to −80 dB, spectral interval aliases into the selected channel bandwidth. Ten such half-band 2-to-1 downsample operations results in 10 such aliased bands folding into the selected channel bandwidth. On the other hand, when the CIC filter performs a 256-to-1 downsample 255 cleared, to −80 dB, spectral intervals alias into the selected channel bandwidth. The two half-band 2-to-1 downsample filters following the CIC fold two additional cleared spectral intervals into the selected channel bandwidth. Thus, the number of low level alias regions that fold into the selected band is 257 for the CIC filter chain and is nine for the half-band filter chain. It is no wonder, that the two filter processes designed for the same minimum attenuation levels lead to different constellation spreading factors. The CIC band stop band edges are the major source of composite alias energy that fold into the final baseband spectral region.

9 Conclusions

In this paper, we presented and compared two efficient solutions for selecting and downconverting a single channel from a large collection of arbitrarily spaced narrow bandwidth channels. These options can be used in portable software radio devices in which the issue of reducing the power consumption is fundamental. Both proposed solutions consist of a 10-stage tree of half-band filters. In one option, the quadrature downconversion is performed as we enter the tree and in the other option it is performed as we leave the tree. The first uses multiple stages of half-band filters with increased number of taps as we move further into the tree. This works out to a nonproblem because as the filter becomes longer, the sample rate is decreased due to the 2-to-1 downsampling per stage.

The second option also uses half-band filters. This option aliases the selected band about the unit circle and tracks its position into one of the four versions of the half-band filter for the next cycle of filter and downsample. A complex heterodyne and a low-pass filter complete the base-band shifting of the selected bandwidth at the end of the processing chain.

We finally noted that the cascade of half-band filters exhibits a smaller spreading about their nominal constellation points than does the CIC filter chain. For the example examined in this paper, the cascade half-band filter chain exhibited a 2.4 dB advantage over the CIC filter chain. We have to thank an anonymous reviewer for suggesting that we examine the relative performance of the two-filter options.

As a suggestion for further research and development, we would like to point out that the dual structure of the channelizer tree presented in this paper can be designed to up convert signal spectra in a portable software radio transmitter, warranting minimum workload and power consumption [5, 6].

References

1. Harris FJ (2004) Multirate signal processing for communication systems. Prentice Hall, Englewood Cliffs
2. Vaidyanathan PP (1993) Multirate systems and filter banks. Prentice-Hall, Englewood Cliffs
3. Harris F, Lowdermilk W (2010) software defined radio: tutorial. IEEE Instrumentation and Measurement Magazine
4. Harris F, Dick C, Rice M (2003) Digital receivers and transmitters using polyphase filter banks for wireless communications. Microwave Theo Tech 51(No. 4):1395–1412, April
5. Harris F, Dick C, Chen X, Venosa E (2011) Wideband 160 channel polyphase filter bank cable TV channelizer. IET Sig Process 5(Issue 3):325–332, June
6. Venosa E, Chen X, Harris F (2010) Polyphase analysis filter bank down converts unequal channel bandwidths with arbitrary center frequencies. SDR Conference, Washington, 30 Nov–3 Dec