

On the Most Efficient M-Path Recursive Filter Structures and User Friendly Algorithms To Compute Their Coefficients

Kartik Nagappa
Qualcomm
kartikn@qualcomm.com

fred harris
San Diego State University
fred.harris@sdsu.edu

ABSTRACT

The standard design procedure for the M-path polyphase filter partitions the impulse response of a prototype low pass FIR filter. The resulting M-path filter, the core of multirate rate filters and filter banks, exhibit the desirable attributes of periodically time varying filters. These attributes being, reduced computational burden for a given filtering task, and simultaneous access to multiple Nyquist zones. While we can not similarly partition the impulse response of a conventional IIR filter, we can design recursive filters as a sum of sub filters with transfer functions of the form $H_n(Z^M)$ thus embedding the desired M-path partition in the design process. In addition to the M-to-1 workload reduction obtained from an M-path polyphase filter, the M-path IIR filter offers an additional 4-to-1 to 10-to-1 reduction relative to the comparable M-path FIR filter. Such filters are described in a number of text books but till recently design algorithms were not available to compute the coefficients required to satisfy a wide range of filter specifications. In this paper we describe the organization of these alternate filter structures, illustrate the result obtained from MATLAB based design routines, and compare the performance and workloads of the M-path recursive and non-recursive filter systems.

1. INTRODUCTION

The sacrosanct Nyquist criteria we learned early in our signal processing education is that the sample rate for a signal should exceed the two sided bandwidth of the signal represented by those samples. It follows then, that when digital filter reduces the bandwidth of a signal, it should, as shown in figure 1, also proportionally reduce the sample rate of the now reduced bandwidth time series.




Figure 1. Digital Filter with M-to-1 Bandwidth Reduction and Sample rate Change

The task of simultaneously reducing the bandwidth and sample rate of a time series is performed by a multirate filter. These filters are traditionally implemented as non-recursive filters that, knowing which output samples are

scheduled to be discarded, simply avoids computing them. We note that when output samples are computed at the reduced output rate of $1/M$ -th the input rate the workload per output sample is amortized over M-input samples thus offering the desired M-to-1 workload reduction.

Traditional recursive filters can not similarly reduce their output rate as part of the filtering process. This is because even when the output port does not require the output sample, and will discard it if offered, the recursive filter must still compute it since the recursion in the filter requires it in order to compute the next output sample.

If we are able to design the IIR filter with segments that only require every M-th sample in its recursion then that filter structure will be able to support the M-to-1 down sampling we obtain effortlessly from the partitioned FIR filter. We examine the process of partitioning the prototype FIR filter, shown in (1), to guide us to the structure of the resampled IIR filter.

$$H(Z) = \sum_{n=0}^{N-1} h(n) Z^{-n} \quad (1)$$

In anticipation of the forthcoming M-to-1 down sampling, as shown in (2), we partition the Z-transform of the prototype filter, into a sum of Z-transforms containing polynomials in Z^M .

$$H(Z) = \sum_{r=0}^{M-1} \sum_{n=0}^{\frac{N}{M}-1} h(r+nM) Z^{-(r+nM)} \quad (2)$$

We can reorder (2) to obtain the polyphase structure, shown in (3), which as a sum of polynomials in Z^M is a form that accommodates the processing of every M-th sample.

$$\begin{aligned} H(Z) &= \sum_{r=0}^{M-1} Z^{-r} \sum_{n=0}^{\frac{N}{M}-1} h(r+nM) Z^{-nM} \\ &= \sum_{r=0}^{M-1} Z^{-r} H_r(Z^M) \end{aligned} \quad (3)$$

The partitioned form of this filter is, as shown in figure 2, a tapped delay line filter with sub filters in Z^M replacing coefficient weights. We note that the M-units of delay at the input clock rate to the sub-filter is the same as one unit of delay at the output clock rate hence we can slide the resampler through the sub-filter, replacing the Z^M at the

input rate with Z^1 at the output rate. The resampled filter is shown in figure 3 where the tapped delay line and the resampler set have been replaced with an equivalent commutator performing the resampling at the filter input rather than at its output.




Figure 2. Polyphase Partition of FIR Filter

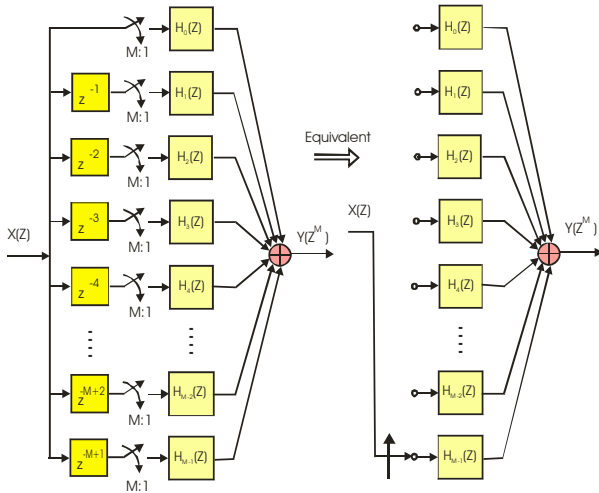


Figure 2. Polyphase IIR Resampling Filters

The polyphase partition of the FIR filter shown in (3) formed filters with Z -transforms written as polynomials in Z^M , a requirement necessary to perform the M -to-1 resampling in the filter. We noted in the introduction, that while it is not generally possible to similarly factor a recursive filter, we can design IIR as polynomials in Z^M for the arms in (3). Filter structures formed in this manner bear little resemblance to recursive filters traditionally implemented as cascade biquadratic (biquad) sections. A useful contender for the arms of the IIR polyphase filter is a cascade of one or more all-pass filters formed as first order and second order polynomials in Z^M , as shown in (4) and in figure 3.

$$H(Z^M) = \frac{1 + a_1 Z^M}{Z^M + a_1} \quad (4)$$

$$G(Z^M) = \frac{1 + b_1 Z^M + b_2 Z^{2M}}{Z^{2M} + b_1 Z^M + b_2}$$

We note that in the implementation of (4) shown in figure (3) a single multiplier coefficient forms both the numerator and denominator of the transfer function and that that the single coefficient simultaneously forms M -poles and a

free set of M -zeros. These zeros reside outside the circle and migrate to the circle when the multiple paths interact through destructive cancellation.




Figure 3. First and Second Order All-Pass Filters in Z^M

Recursive filters defined with denominators formed by polynomials in Z^M have roots residing on vectors aligned with the M -roots of -1 . Figure 4 illustrates the root location alignment with the 8-roots of -1 for one path of an 8-path filter formed by a cascade of three first order, in Z^8 , all-pass filters. Note that each multiply forms 8-poles and 8-zeros and the three sub filters in this path form 24 poles and 24 zeros. If each path contributes a comparable number of poles and zeros the composite filter would contain 192 poles and 192 zeros. Not bad for 24 multiplies!

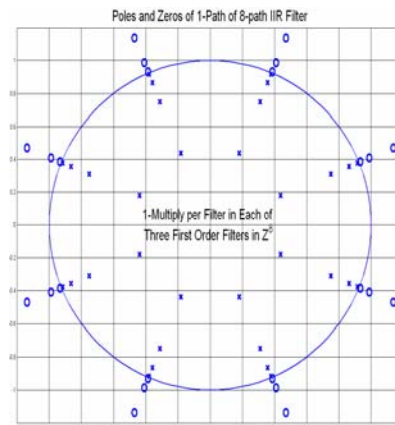


Figure 4. 24-Poles and Reciprocal Zeros of Three First Order All Pass Filters in Z^8

We note that all-pass filters have a steady state gain of unity at all frequencies and are characterized by frequency dependent phase shift. All-pass filters find wide application in phase equalizers that add frequency dependent phase to the phase of arbitrary recursive filters to obtain a composite linear phase profile. We also note that products (or cascades) of all-pass filters are still all-pass filters and

that sums of all-pass filters are not all-pass. We use this latter property to design our filter.

2. M-PATH FILTER PROPERTIES

The phase profiles of the separate paths in the M-path filter are constrained in the following manner. Over the frequency span assigned to the pass band, the phases are nearly identical. The spectral regions corresponding to this condition add coherently when the signals from the M-paths are added and thus exhibit a gain of M which is scaled to unity by multiplying by 1/M. Over the frequency span assigned to the stop band the phases differ by multiples of $(2\pi/M)$. Since the M roots of unity sum to zero, the spectral regions corresponding to this condition add destructively to zero thus forming the desired stop band. Figure 5 illustrates the phase profiles associated with 8-path filters. We see here three profiles corresponding to a linear phase FIR, a linear phase IIR, and a non-linear phase IIR 8-path filter. Here we see the phase profile are identical in the low frequency band and then separate in successive frequency bands to form the stop band.

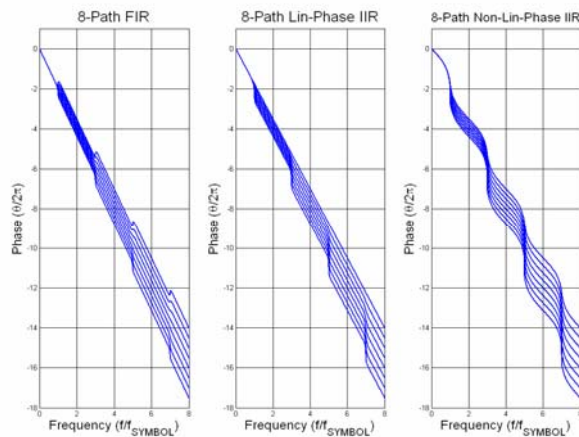


Figure 5. Phase Profiles of Each Path in 8-Path Filter

Figure 6 shows the same phase profile except the profiles have been de-trended by subtracting the average phase (or zero-th path phase) from the profiles. Here we see the constant phase differences in each successive frequency band that support the desired destructive interference. The regions which do not exhibit the constant phase difference are the transitions between successive bands and these transitions are responsible for amplitude transitions in M-path IIR filters.

The linear phase version of the M-path IIR filter is obtained by setting the zero-th path in the set to a delay, the simplest all-pass filter. The design algorithm adjusts the phase of the remaining paths to match the phase slope of this reference path. None of the phase slopes are predetermined in the non-linear phase IIR filter, and this additional degree of freedom offers solutions requiring a small

number of sub filters in each path. The two design options require different filter design algorithms.

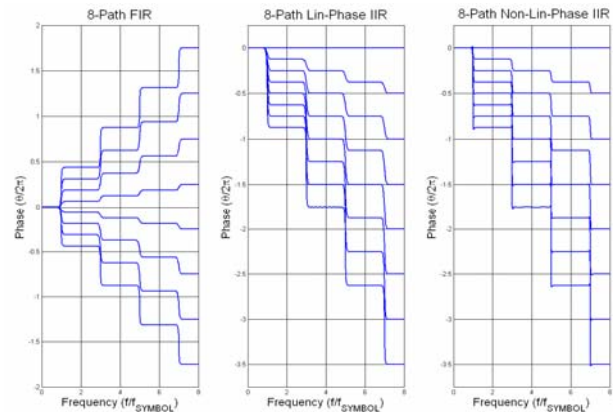


Figure 6. De-Trended Phase Profiles Paths in 8-Path Filter

Since perfect destructive cancellation will be observed as a stop band zero, and since only a finite number of zeros reside in the stop band, the interval between the zeros exhibits stop band ripple with spectral attenuation levels controlled by the filter design parameters. The spectral region between the pass band and stop band is of course the transition bandwidth of the filter.

The angles on the unit circle matching the sectoring lines shown in figure 3 correspond to the phase transition frequencies shown in figures 5 and 6. The frequency response between these transitions represents successive Nyquist zones in an M-path filter. Due to the aligned poles, an M-path recursive filter exhibits a gain transition at the frequency matching the phase transitions while the M-path non-recursive filter, without these aligned poles of course does not. When present, these transition bands are treated as don't care zones similar to the side lobes between filter zeros in the CIC filter.

3. M-PATH FILTER PERFORMANCE

Traditional wisdom in the signal processing community is that recursive filters are always more efficient than a non recursive filter designed to same performance specification. This proves not to be so when we include the option of sample rate change as suggest in figure 1. When we teach filter design with emphasis on system applications, we demonstrate that a resampled FIR filter always requires fewer operations than its equivalent IIR. Another practical consideration is that most DSP architectures are optimized to perform inner products, and are capable of performing pipelined multiply accumulate operations more efficiently and in less time than a competing IIR implementation. We then claim that there is never a need to implement IIR filters more sophisticated than a one

pole filter, the leaky integrator. This proclamation is based on the FIR filter's ability to avoid computing output samples that will be discarded in a subsequent down sampling operation, an option that is not available to the IIR filter. This comparison is valid for an M-to-1 down sampling filter as well as for a 1-to-M up sampling filter.

When we add the M-path IIR filter as an option in the comparison process, we find the results are different. In fact, if computational burden is the only criteria, the M-path IIR filter outperforms the M-path FIR filter. To illustrate the relative workload and system performance we designed a set of 8-path filters to reduce the sample rate, by a factor of 8-to-1, of a time series collected by a DAC that collected data at 8-times the desired final sample rate. It is common to control the cost of analog anti-aliasing filters by collecting data at an over sampled rate and then reducing the bandwidth and sample digitally. The digital filter specifications for our example include a normalized frequency pass band from 0-to-0.9 with less than 0.1 dB pass band ripple with stop bands ± 0.9 centered at 2, 4, 6, and 8, the multiples of the input sample rate 2. Stop band attenuation must exceed 72 dB.

Figure 7 presents the impulse response and frequency response of a FIR filter meeting these system requirements. Shown in the spectral plot is the attenuation boundary demarcations which also bound the don't-care bands distributed in the stop band region. The FIR filter does not use the don't-care bands. The prototype filter is seen to be of length 128 taps which when partitioned into an 8-path filter lead to a decimation filter requiring 16 multiplies per input point.

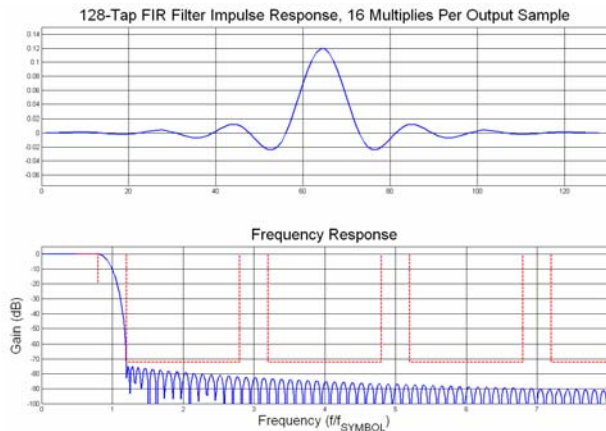


Figure 7. Time and Frequency Response: FIR Filter

Figure 8 presents the impulse response and frequency response of an 8-path linear phase IIR filter meeting the same system requirements. Here we see the don't-care bands between the demarcated attenuation boundaries are now occupied by the transition bands we discussed earlier. Seven of the paths in the 8-path filter contain two

first order filters in Z^8 in cascade with three second order filters in Z^8 , requiring 8 multiplies per path for a total of 56 multiplies in the entire filter. When amortized over the 8-input samples, the workload per sample is seen to be 7 multiplies per sample. This workload is less than half that of the 8-path FIR filter. We keep in mind that this filter is recursive and does not take advantage of the efficient multiply-accumulate option available to the FIR filter.




Figure 8. Time and Frequency Response: Prototypic 8-Path Linear Phase IIR Filter

Figure 9 presents the impulse response and frequency response of an 8-path non-linear phase IIR filter meeting the same system requirements. Here again we see the don't-care bands between the demarcated attenuation boundaries are now occupied by the transition bands we discussed earlier. The non-linear phase 8-path IIR structure contains four first order filters in Z^8 in the first four paths and three first order filters in Z^8 in the next four paths. The total number of multiplies in the filter is seen to be 28 and when amortized over the 8-input samples, the workload per sample is seen to be 3.5 multiplies per sample. This workload is again less than half that of the 8-path linear phase IIR filter. Again, we note that this filter is recursive and does not take advantage of the efficient multiply-accumulate option available to the FIR filter.

We commented earlier that the efficiency of the M-path IIR filter is related to the fact that each multiply spawns M poles and M zero. The number of roots formed by the aggregate of M paths containing multiple polynomials in Z^M is quite large and is orders of magnitude greater than the number of roots we normally consider in conventional IIR filter structures. By way of comparison, an inverse Tchebyshev filter satisfying the same specifications as the FIR filter shown in figure 7 requires 11 poles and zeros with an additional 8 poles and zeros if phase equalized. The 11-pole filter implemented by biquad sections requires 28 multiplies to form the roots and appropriate distributed scale factors. An elliptic filter meeting the

filter specification requires 7 poles and zeros and when implemented by biquad sections requires 18 multiplies. Note this supports our earlier contention that a resampled FIR filter (16-M) out performs the comparable IIR filter (18-M).

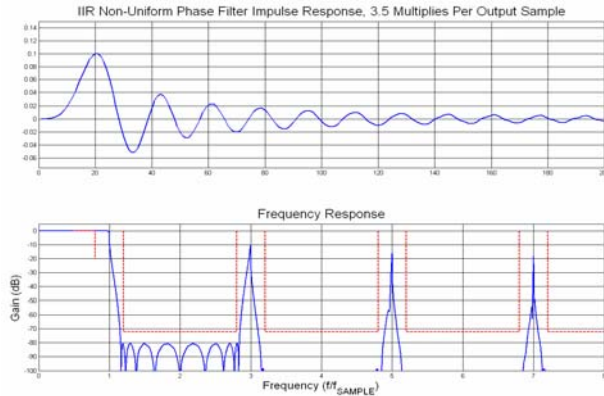


Figure 9. Time and Frequency Response: Prototype 8-Path Non-Linear Phase IIR Filter

As a reference for comparison, figure 10 presents the pole zero diagram for the 128 tap FIR filter. As expected all the poles reside at the origin and the zeros reside on the unit circle as stop band zeros or if not on the circle appear off the circle in reciprocal pairs. We see that 12 zeros in the pass band region form reciprocal pairs with 6 non-minimum phase zeros outside the circle, a characteristic of linear phase filters. The remaining 115 zeros are stop band zeros.

Figure 11 shows the pole zero diagram of the 8-path, linear phase IIR filter. This filter required 64 delays in the reference path, path-0, and 2 first order filters in Z^8 in cascade with 3 second order filters in Z^8 for the remaining 7 paths. Accounting for the additional delays inserted in each path we find that the filter contains a total of 520 poles and 514 zeros. Of this number, 119 zeros are stop band zeros, 7 are non-minimum phase zeros, and the remaining zeros act to shield the influence of the distributed poles on the stop band zeros. We can also identify the poles aligned with the 8-roots of -1 near the unit circle responsible for the phase and gain transition bands.

Figure 12 shows the pole zero diagram of the 8-path, non-linear phase IIR filter. This filter required 4 first order filters in Z^8 in the first four paths and 3 first order filters in Z^8 in the remaining four paths. Accounting for the additional delays inserted in each path we find that the filter contains a total of 232 poles and 226 zeros. Since the filter only uses first order filters in Z^8 we note that all poles are aligned with the 8-roots of -1. Of the 226 zeros, 63 are

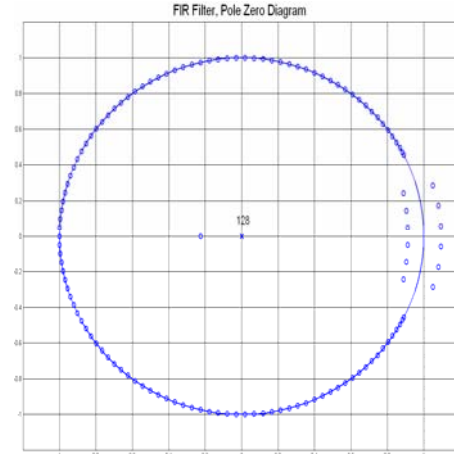


Figure 10. Pole-Zero Plot of 128 Tap FIR Filter




Figure 11. Pole-Zero Plot for Prototype 8-Path Linear Phase IIR Filter

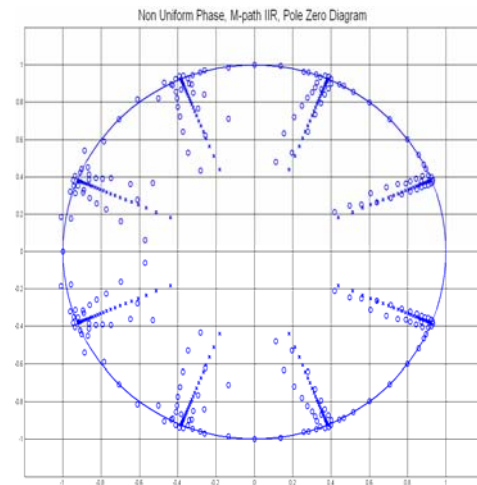


Figure 12. Pole-Zero Plot for Prototype 8-Path IIR

stop band zeros, with the remaining zeros buffering the influence of the distributed poles and the stop band zeros.

4. NYQUIST ZONE FILTERING

The 1-to-M up sampling or the M-to-1 down sampling embedded in the M-path polyphase filter enables the filter to operate in any Nyquist zone of the resampling process. As a result of the resampling process that occurs prior to the filtering process, all the Nyquist bands alias to base band. Due to their different pre-alias center frequencies, the aliases have distinct phase profiles thus can be separated by their unique phase profiles. Earlier we demonstrated the constant phase offset between each path of the 8-path filter and the different values of phase offset in the different Nyquist bands. These phase profiles are identical to the phase profiles of the signal spectra in the aliased Nyquist zones. We can extract the signal from any of these zones, even if they have aliased to base band, by setting the phase difference in that zone to zero. We accomplish this by applying complex phase rotations to the signal outputs from each path. The zone that exhibits zero phase difference in the phase rotated polyphase filter is the one that survives the summation of the path sums. This process is equivalent to a cascade of a quadrature down conversion, a pair of low pass filters, and a down sampler but operated in the opposite order. By reversing the order, the filter is real since the data does not become complex till after the filter as opposed to becoming complex prior to the filter. This is shown in figure 13.




Figure 13. Polyphase Partition and Phase Rotators to Extract Alias from k-th Nyquist Zone

Suppose we want to extract the signal from the selected Nyquist zone but we do not want to translate the signal to base band nor reduce the sample rate. It would seem that the workload could not be reduced because we have been denied access to the resampling option. The resampling option is so efficient we can embed it in the solution. What we do is down sample and extract the alias at the reduced sample rate and then up sample with the phase rotators to reinsert the spectra at its initial (or different) Nyquist zone. This option is shown in figure 14. If we use the linear-phase IIR filter to accomplish this, the work-

load is 9 multiplies per input sample and 9 multiplies per output sample and since the input and output rate are the same, the workload is 18 ops per input/output sample. A standard IIR elliptic band-pass filter meeting these specifications has 16-poles and requires 40 multiplies while a comparable FIR filter requires 128 multiplies. The computational cost of down sampling, alias extracting, alias assignment and up sampling is seen to be smaller the single band-pass filter.

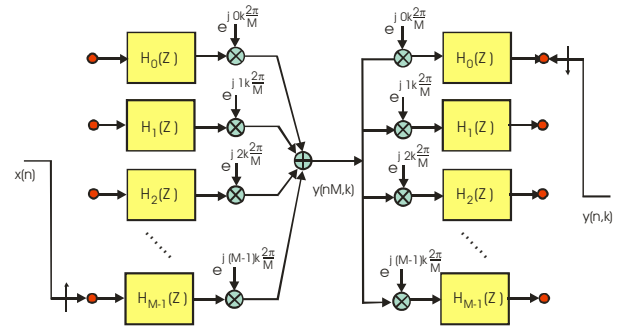


Figure 14. Cascade Polyphase Down Sampler and Up-Sampler to Extract k-th Nyquist Zone

5: CONCLUSIONS

We have discussed and reviewed the structure of digital recursive filters that are capable of being partitioned into M-path polyphase filters in a manner similar to the partition of the finite duration impulse response of a non recursive filter. This structure is the sum of products of all-pass, first order and second order polynomials in Z^M , filters. We demonstrated some of the unusual properties of this class of filters including the phase profiles of the separate paths, the presence of don't care spectral bands, and the existence of unusually large number of roots in the composite filter. We also demonstrated the significant reduction in computation load that can be realized using polyphase IIR filters

6: REFERENCES

- [1]. Markku Renfors and Tapio Saramäki, "Recursive Nth-Band Digital Filters – Part I: Design and Properties", IEEE Trans. On Circuits and Systems, Vol. CAS-34, No. 1, January 1987.
- [2]. fred harris, "Multirate Signal Processing for Communication Systems", Prentice-Hall, 2004, Ch. 10, pp. 258-324.
- [3]. fred harris and Steve Scott, "Flexibility, Performance, and Implementation Advantages of Recursive, Linear and Non linear Phase, Polyphase Filters in Transmitters and Receivers", SDR-04, Phoenix, AZ, 2004.

