ARCHITECTURE FOR MINIMUM POWER RECURSIVE DIGITAL FILTERS

fred harris (San Diego Sate University, San Diego, CA, fred.harris@sdsu.edu)
Gregory Smith (Signum Concepts, San Diego, CA, gregg.smith@Cubic.com)
Richard Jeckel (Cubic Corporation, San Diego, CA, Richard.Jeckel@Cubic.com)

1.ABSTRACT

Software defined radios require filter architectures that are flexible, easily reconfigurable, and boast low power consumption. The need for low power consumption leads us to multirate signal processing in which one or more sample rate changes are embedded in the filtering process to permit operation at the lowest rate commensurate with the signal bandwidth. It is a simple matter to embed the sample rate change in a FIR filter by not computing the samples destined to be discarded. This is an option not generally available to recursive filters. This paper presents a recursive filter architecture that does permit this sample skipping. We present a typical filtering task implemented by various design approaches and compare their computational workloads. The workload for the filter architecture presented here is typically one-fourth to one-sixth of conventional architectures.

2. INTRODUCTION

Multirate filtering has become the standard response to power efficient digital signal processing tasks. The wisdom that motivates us to pursue multirate processing comes from the design equation that relates filter length to filter performance. The parameters that define filter performance are shown in figure 1 and the relationship that define the filter length "N" in terms of these parameters is shown in equation 1. Equation 2 presents first order approximations to equation 1 for FIR filters designed by windowing and by the Remez algorithm. N, the filter length indicates the number of multiplies and add required to compute each has output sample. In this light, N has units of ops/output, and we are motivated to minimize this parameter.

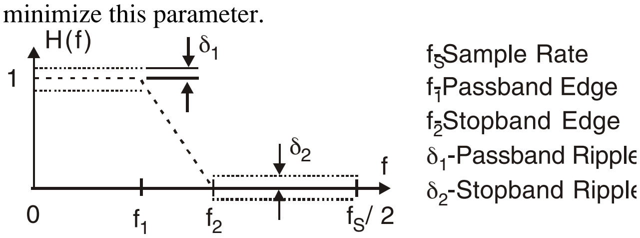

Figure 1. Parameters of Sampled Data Low-Pass Filter

(1)

(2)

The important concept contained in equation 2 is that filter length will increase if you require greater out-of band attenuation or a narrower transition bandwidth. These parameters are specified to satisfy system performance requirements and cannot be changed to obtain reduced length filters. The one free parameter we are still free to control is the sample rate f_S. If we reduce the sample rate by a factor of M so that the filter operates at f_S/M, the workload is reduced by the same factor M. The block diagram of the filter and resample process is shown in figure 2 and the modification to equation 2 to reflect the sample rate change is shown in equation 3. In reality, we interchange the order of filtering and resampling, and only compute the output samples required by the resampling switch. We manage this by computing one output sample for every M input samples. The filter architecture is often modified to be an M-path filter to accommodate this process reordering, and the modified structure is called a polyphase filter. In this modified form, each path of the filter performs N/M operations per input sample performing a portion of the workload required to obtain the N operations per output sample.

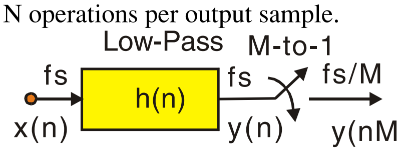

Figure 2 Resampling Low Pass Filter

(3)

We often perform the resampling and filtering task as a cascade of two filters, one that performs an initial low-cost reduction in bandwidth and sample rate which is followed by a clean up filter that performs a final reduction in bandwidth and sample rate as well as corrects distortion caused by the low cost pre-filter. This structure is shown in figure 3. The input filter is designed with a wide transition bandwidth to obtain a reduced filter length N_1. The second filter now designed to operate at the reduced sample rate is implemented with a reduced filter length N2.

If we have a requirement that the output sample rate is to be same as the input sample rate, we modify the processing stream to return the sample rate to the original by adding an up sampling filter. The up sampling filter is the dual of the down sampling filter and it too is a polyphase filter but as a dual filter, the workload is N/M ops/output sample as opposed to N/M ops/input sample. This option is shown in figure 4.

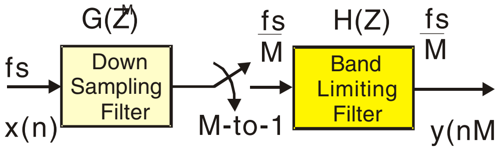

Figure 3 Resampling as a Cascade of Two Filters

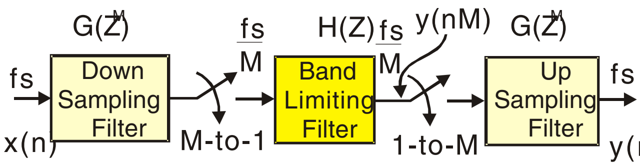

Figure 4 Resampling as a Cascade of Three Filters

We may have a sense of discomfort questioning whether three filters can offer a reduced workload relative to the single filter it is replacing. In response to this concern a number of investigators have suggested another standard option to obtain a reduced workload. This option is know as the Interpolated Filter or IFIR technique in which we leave the sample rate of the data fixed but design the filter at a reduced sample rate and then up sample the filter by a zero-packing process. The zero packed filter exhibits multiple pass bands, which are rejected by the second interpolating filter. The interpolating filter has a wide transition bandwidth hence is of low order. The IFIR structure is shown in figure 5. Note that the two filters shown in figure 5 are the same filters shown in figure 3 but operating in the reverse order at different sample rates.

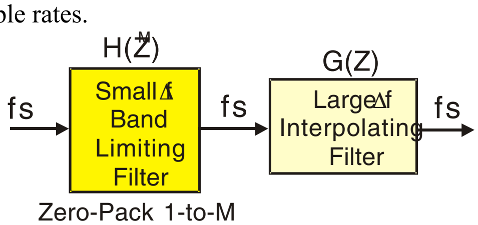

Figure 5. IFIR Filter: Zero Packed 1-to-M

3. WORKLOAD COMPARISON

We now examine a specific, but typical, signal-processing task and compare the workload required for each of the standard implementation options. This is followed by a presentation of an alternate, non-standard recursive filter structure that is then compared to the standard options. Specifically we consider the implementation of a low pass digital filter with the following specifications:

Sample Rate:	100 kHz		
Pass Band:	0-to-5 kHz	In Band Ripple	0.1 dB
Stop Band:	7-to 50 kHz	Stop Band Attn	60 dB

A single stage FIR filter that meets these specifications requires 150 taps. Figure 5 presents the impulse response and frequency response of the single stage filter.

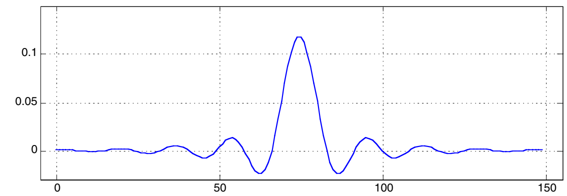

Figure 6. Impulse and Frequency Response of 150 Tap Filter

The specifications for an interpolated 2-stage filter designed to operate the input stage at one-fifth of the input rate, or 20 kHz followed by the interpolating filter operating at the input rate of 100 kHz are listed below. As expected the first stage filter length is one-fifth the length of the full sample rate prototype filter. The interpolating filter has a length of 40 taps so the arithmetic workload for the cascade is 30+40 or 70 ops per input as opposed to the 150 ops per input for the single stage prototype.

	First Stage	Second Stage
Filter Type	Low Pass	Low Pass
Sample Rate	20 kHz	100 kHz
Pass Band	0-to-5 kHz	0-to-5 kHz
Stop Band	7-to 10 kHz	15-to-50 kHz
In Band Ripple	0.08 dB	0.02 dB
Stop Band Attn	60 dB	60 dB
Filter Length	30 taps	40 taps

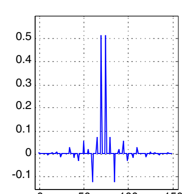

Figure 7. Impulse and Frequency Response of IFIR Filters

Figure 7 presents the impulse response and frequency response of the two filter stages as well as the cascade of the two stages.

As a final comparison we can form a cascade of three filters that perform 5-to-1 down sampling, band limiting, and 1-to-5 up sampling of the form shown in figure 4. The input and output stages of this cascade are polyphase partitions of the interpolator in the example just described while the band liming filter is the non zero packed version of the first stage filter from the same example. Figure 8 presents one of the five impulse responses that can be observed at the filter output plus the frequency response of the three stages and the composite response of the three stages. The frequency responses have all been generated at the final output sample rate.

When this third option filter operates, for every 5-input points, the workload performed by the three stages are 40 ops, 30 ops, and 40 ops respectively for a total of 110 ops per 5 input samples. The workload per sample for this option is 22 ops per input sample. Interestingly, this option exhibits the smallest arithmetic workload even though it exhibits the longest composite impulse response. The length of the impulse response can be determined in the following manner. An applied impulse outputs an 8-point sequence from the first filter corresponding to one of the 5-legs of the input polyphase filter. This sequence then passes through the 30-tap filter to output a sequence of length 37 points. This 37-tap sequence in turn moves through the five legs of the output polyphase filter to output a sequence of length 224 points (37*5+40-1). Table 1 compares the arithmetic workload and the equivalent filter length of the three filter options examined in this example.

Option	Filter Coefficients	Arithmetic Workload	Impulse Response Length
Single Stage	150 Taps	150 ops/input	150 Taps
Interpolated Filter	30, 40 Taps	70 ops/Input	189 Taps
5-to-1 and 1-to-5 Resampling Filter	40, 30, 40 Taps	22 ops/input	224 Taps

Table 1. Comparison of Structures: Coefficient Lengths, Arithmetic Workload, & Composite Filter Impulse Response

In this comparison we note an observable truism about signal processing. It is this: To obtain the lowest processing burden, *a signal-processing task should always be conducted at the lowest possible sample rate* commensurate with the signal bandwidth!

4. RECURSIVE FILTER OPTIONS

For completeness, we now examine conventional recursive filters to determine the workload required for the filtering task under discussion. From standard recursive filter design routines we determine that a 7-th order elliptic filter will satisfy the filter specifications. At first glance, it appears that a recursive filter with 7-poles is a more efficient option than the non-recursive filter with 150-taps. We must be careful here, comparing poles of a recursive filter to coefficients of a non-recursive filter does not make sense. This is the apples and oranges problem. What we should compare is ops per data sample. In the FIR filter case, we found the best option to be 22-ops/data sample. We now determine the corresponding workload of the recursive filter.

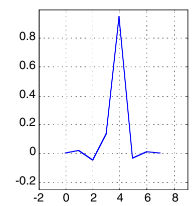

Figure 8. Impulse Responses Cascade Resampling Filters and Composite Response Along with their Frequency Responses

Figure 9 shows a standard implementation of the 7-th order filter as a cascade of three 2-nd order canonic filters and one 1-st order filter. The second order filters are often called biquads for biquadratics. Each 2-nd order filter requires five multiplies, two for the feedback path and three for the scaled feed forward path while the 1-st order requires three multiples, one for the feedback and two for the scaled feedforward The 7-th order filter requires a total of 18 multiplies per input-output data pair. At the moment, this filter requires the smallest workload but if we had to include a phase equalizer, the resampled FIR would be the better option.

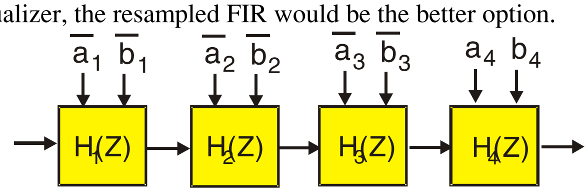

Figure 9 Canonic Filters for 7-th Order Elliptic Filter

Figure 10 presents the impulse response and the frequency response of the standard 7-th order elliptic filter. We now introduce and compare other designs that offer reduced computation.

5. TWO PATH RECURSIVE POLYPHASE FILTERS

An alternate filter structure to the cascade biquadratic filter is the two-path polyphase filter structure. This structure is shown in figure 11. There are two elementary building blocks in the structure and they are the 2-nd order and 1-st order recursive all-pass filters shown in figure 12. These filters have unity gain at all frequencies and only contribute a frequency dependent phase shift to the gain of each path. Note here that each coefficient is in the feed forward and the feedback path of the filter and hence half the number of coefficients are required to form the filter poles and zeros

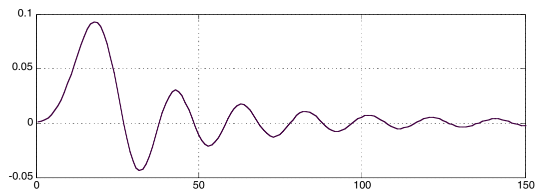

Figure 10 Responses of 7-th Order Recursive Filter

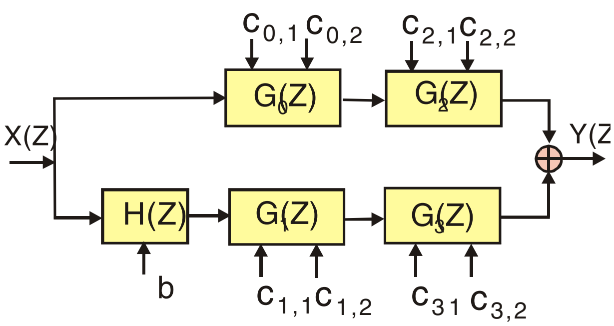

Figure 11. Two-Path Recursive Polyphase Filter

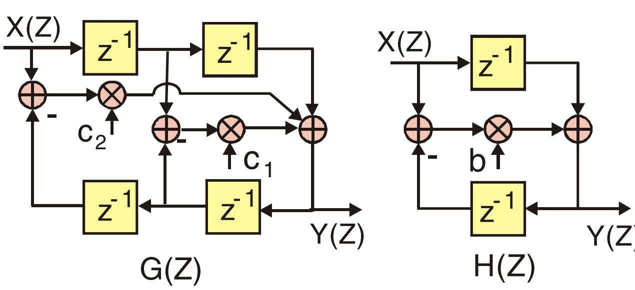

Figure 12. Second and First Order Recursive All-Pass Filters

In the two path filter architecture we arrange for the phase of the all-pass filters in the two paths to match in the frequency span we define as the pass band while simultaneously having the phase of the two paths differ by 180 in the frequency span we define as the stop band. The filter stop band is formed when the signals from the two paths are added and experience the destructive cancellation of the spectral terms that differ by 180 degrees. We wrote a MATLAB script file called *Tony_des* to compute the filter coefficients for the architecture. This program is named after Tony Constantinides of Imperial College who first introduced us to this class of filters. Nine coefficients are required to satisfy the specifications of the problem we are examining. The impulse response, the phase response of the two paths in the filter, and the frequency response of this filter are presented in figure 13. Here we clearly see that the filter stop band corresponds to the frequency span in which the two paths differ by 180 degrees and thus enable destructive cancellation when the signals at the output of the two paths are summed.

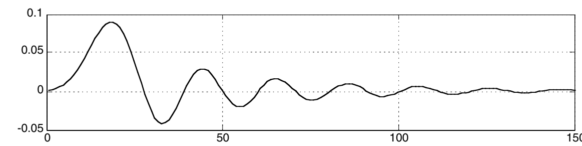

Figure 13 Impulse Response, Path Phase Profiles, and Frequency Response of Recursive Two Path Filter

6. RESAMPLING RECURSIVE ALL-PASS FILTERS

We commented earlier that processing occurs most efficiently at the lowest sample rate commensurate with the signal bandwidth. As efficient as these two-path recursive all pass filters are, there is additional processing savings to be had by using the filters in a resampling mode. Classic recursive filters cannot be resampled because recursive filters need the previous output sample in order to form the next output sample. Thus even if we have no interest in the filter output, the general recursive filter cannot share our disinterest. A pleasant surprise is that the two-path recursive filter can be resampled when it is used as a half-band filter. We thus have the option to cascade two half-band two-path filters and lower the sample rate by 4, perform the filtering at the reduced sample rate, and then return the sample rate to the original by up sampling with two more half band-filters. This is in accord with the structure presented in figure 4. Using *Tony_des* we determined that the first half band filter required two coefficients, that the second half band filter required 3-coefficients and the center bandwidth limiting filter required 7 coefficients. When cast in their appropriate locations and operating as half-band down sampling filters at the input and as half-band up sampling at the output we obtain the structure shown in figure 14.

Note that the input filter has the interesting property that it only requires a single multiple per input sample. The second filter performs 3 multiples per pair of inputs delivered to it which amortizes to 3/2 multiplies per delivered input, but it sees samples at half the input rate which we acknowledge by stating that the second filter performs 3/4 of a multiple per input sample. The middle filter performs 7 multiplies per data sample delivered to it and reflecting this workload to the 4-times higher input rate we arrive at the observation that the middle filter requires 7/4 multiples per input sample.

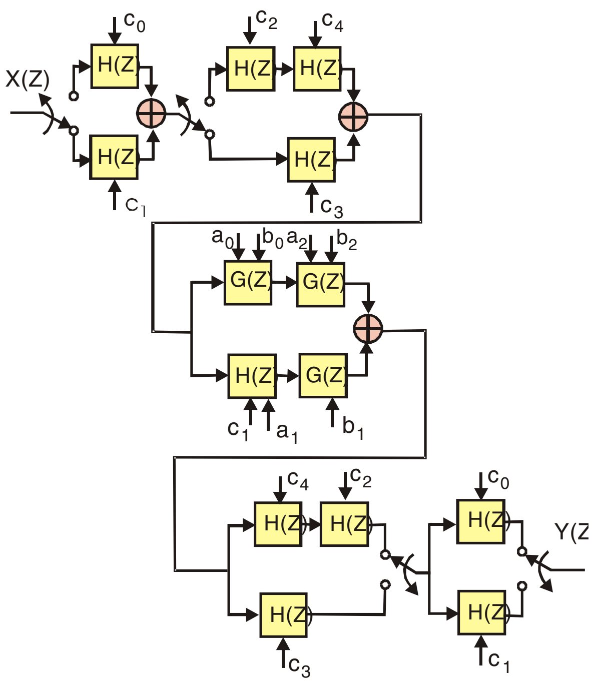

Figure 14. Cascade Half-Band Down Sampling, Filtering and Half Band Up Sampling Two-path Recursive All-Pass filters

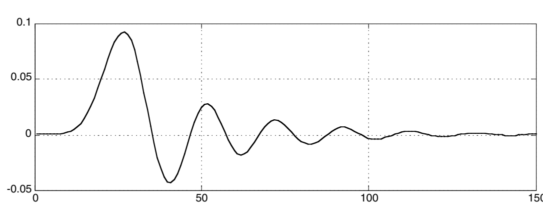

Figure 15. Impulse and Frequency Response of Resampled Filter

Summing the workload over the 5 filters we see the total workload per input sample is [1 + 3/4 + 7/4 + 3/4 + 1] or 21/4, which is 5.25 multiplies per input sample.

Figure 15 presents the impulse response as well as the frequency response of the composite and of the separate half band filters. The spectra of the half band filters are presented at the common input sample rate. Table 2 presents a comparison of the complexity and of the workload for requirements for the three-recursive filters we have just examined.

Option	Filter Coefficients	Arithmetic Workload
7-Pole Biquadratic Elliptic Filter	7-Fdbk,11-Fdfwrd	18 ops/input
9-Pole Two Path	9-Fdbk/Fdfwrd	9 ops/input
2-to-1, 2-to-1 and 1-to-2, 1-to-2 Resampling IIR	2,3,7,3,2 Coefficients	5.25 ops/input

Table 2 Comparison of Structures: Coefficients and Workload

Figure 16 is a bar graph presenting the workload corresponding to the six filters, we have examined in this paper. On the high end we have the single stage FIR at 150 ops/input while at the low end we have the cascade of resampled recursive two-path filters at 5.25 ops/input.

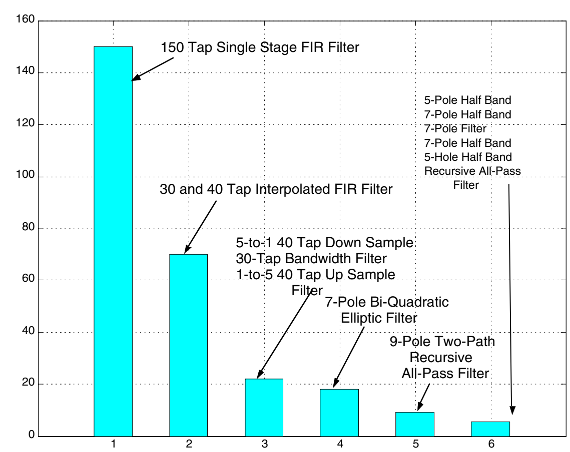

Figure 16 Bar Graph presentation of Workload for Three FIR Filters and Three IIR Filters

A final comment here is that we have compared recursive and non-recursive filters without regard for their group delay distortion. The FIR filter offers linear phase response in all its variants while the IIR filters we examined were all characterized by severe group delay distortion. The IIR filters could have had a phase equalizer appended to the bandwidth-limiting filter but that would have doubled their workload, still offering a respectable option. In another option, the two IIR two-path half-band filters could have been implemented as equiripple group delay filters by setting the top path to delay only. In this case, the lower path becomes an equiripple approximant to a linear phase filter and the half band filter would inherit the linear phase property.

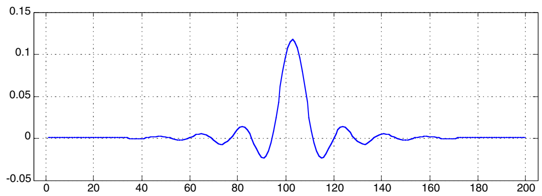

Figure 17 Impulse and Frequency Response of Linear Phase Recursive Outer Filters with FIR Inner Filter

Figure 17 presents the impulse response obtained from a filter with two half-band two-path linear-phase recursive polyphase outer filters, requiring 2 and 3 weights, and a 36-tap FIR inner filter. The workload per input sample point is [4 + 3 + 36 + 4 + 3]/4 or 12.5 ops/sample. This linear phase hybrid-resampling cascade requires approximately half the workload of the FIR resampling cascade. The interested reader is invited to contact the authors for additional material on linear phase versions of the two-path polyphase recursive filter.

6. CONCLUSIONS

We have presented and compared a number of methods to realize a bandwidth limiting digital filter. The most straightforward approach is the single stage non-recursive filter, which is also the most expensive in terms of workload, requiring 150-ops/input sample. By clever manipulation of the sample rate through the use of resampling filters we were able to pull the workload down to 22 ops/input. We then shifted approaches and examined the straightforward recursive filter. This resulted in a 7-pole elliptic filter, which required 18-ops/input sample when we included feedback, feed forward, and inter-stage scaling. We then introduced the two-path all-pass filter structure in which each multiply formed a system pole and a reciprocal, outside the unit circle, zero. Addition of the two all-pass transfer function moved the zeros from their original locations positions to the unit circle. The zero shift is accounted for by destructive cancellation of the signals from the two filter paths. Since a single coefficient and its associated multiplication forms both a pole and a zero, the two path filter is always more efficient than the equivalent performing cascade biquadratic filter by a factor of 2-to-1. Finally we accessed another nearly 2-to-1 reduction in workload by invoking sample rate changes available through very efficient half-band filters formed by the two-path half-band recursive filter. The recursive filter options required 18, 9,and 5.25 ops per input sample to effect the desired bandwidth reduction while the linear phase-hybrid, part recursive, part non-recursive required 12.5 ops per input sample. We noted that the hybrid required about half the workload of the best linear phase all FIR resample solution.

7. REFERENCES

P.P. Vaidyanathan, "Multirate Systems and Filter Banks", Prentice-Hall, 1993.

R.A. Valenzuela and A.G. Constantinides, "Digital Signal Processing Schemes for Efficient Interpolation and Decimation", IEE Proceedings, Dec 1983

M. Renfors and T. Saramaki, "Recursive N-th Band Digital Filters, Parts I an II", IEEE Trans. CAS, Vol. 34, pp. 24-51, January 1987.

fred harris, Bob Caulfied, and Bill McKnight, "Use of All-pass Networks to Tune the Center Frequency of Sigma Delta Modulators", 27th Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA., 31 October to 3 November 1993

fred harris and Eric Brooking, "A Versatile Parametric Filter Using an Imbedded All-Pass Sub-Filter to Independently Adjust Bandwidth, Center Frequency, and Boost or Cut, 95th Audio Engineering Society (AES) Convention, New York, New York, 7-10 October 1993

fred harris, Itzhak Gurantz, and Shimon Tzukerman "Digital (T/2) Low pass Nyquist Filter Using Recursive All-pass Polyphase Resampling Filter for Signaling Rates 10 kHz to 10 MHz", 26th Annual Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA., 26–28 October 1992

fred harris "On the Design and Performance of Efficient and Novel Filter Structures Using Polyphase Recursive All-Pass Filters", Keynote Presentation ISSPA-92 (International Symposium on Signal Processing) Gold Coast, Australia 16–21 August 1992

fred harris, Maximilien d'Oreye de Lantramange and A.G. Constantinides, "Design and Implementation of Efficient Resampling Filters Using Polyphase Recursive All-Pass Filters", 25th Annual Asilomar Conference on Signals, Systems, and Computer, Pacific Grove, CA., 4–6 November 1991

fred harris, "Multirate Signal Processing for Communication Systems: *Current Practice and Next Generation Techniques*", Prentice Hall, 2004.