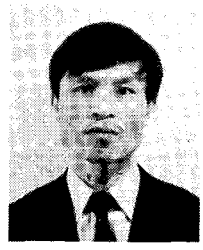
the forced negative-resistance oscillator," *IEEE Trans. Circuits Syst.*, vol. CAS-28, pp. 217–224, 1981.
[16] C. Hayashi and H. Kawakami, "Bifurcations and chaotic states of the solutions of Duffing–Van der Pol equation," in *Proc. 26th Int. Wiss. Koll.*, TH Ilmenau, vol. A-1, pp. 15–18, 1981.
[17] M. J. Feigenbaum, "Qualitative universality for a class of nonlincar transformations," *J. Stat. Phys.*, vol. 19, pp. 25–52, 1978.
[18] H. Kawakami and Y. Katsuta, "The Hopf bifurcation and chaotic states of solutions in a third order equation of Duffing's type," *Trans. IECE Japan*, vol. J64-A, pp. 940–947, 1981 (in Japanese).
[19] E. I. Jury, *Inner and Stability of Dynamical Systems*. New York: Wiley, 1974.



Hiroshi Kawakami (M'81) was born in Tokushima, Japan, on December 6, 1941. He received the B.Eng. degree from Tokushima University, Tokushima, Japan, in 1964, the M.Eng. and Dr.Eng. degrees from Kyoto University, Kyoto, Japan, in 1966 and 1974, respectively, all in electrical engineering.

Presently, he is Associate Professor of Electronic Engineering, Tokushima University, Tokushima, Japan. His interest is in qualitative properties of nonlinear circuits.

Some Efficient Digital Prefilter Structures

JOHN W. ADAMS, MEMBER, IEEE, AND ALAN N. WILLSON, JR., FELLOW, IEEE

***Abstract*—A new approach to the design of low complexity FIR digital filters was presented in [1]. The essence of the method is to separate the design problem into two parts: the realization of an efficient prefilter and the design of the corresponding equalizer. This separation allows the filter designer to focus on the computational complexity issue within the simplified context of the prefilter network.**

In this paper, we consider the use of various digital prefilter structures. The number of possibilities for digital prefilters is unlimited, so that our treatment cannot be exhaustive. However, our study does provide a quantitative sense of the level of reduction in computational complexity that can be achieved by the use of an appropriately designed prefilter.

I. INTRODUCTION

A NOVEL approach to the design of efficient FIR digital filters was presented in [1]. The essence of the proposed method is to decompose the design problem into two parts:

(1) Extract an efficient "prefilter" design from the specifications. The prefilter should have the best possible frequency response with a minimal number of multipliers and adders.
(2) Cascade the prefilter with an "equalizer" to achieve the desired overall result.

There exists a very simple and efficient low-pass prefilter structure. It has only two adders and no multipliers at all, regardless of the filter length. It is referred to as the RRS (recursive running sum). Fig. 1(a) shows the RRS structure which performs only two additions to obtain each output sample. Fig. 1(b) shows the mathematically equivalent direct-form implementation. The RRS is just an FIR digital filter with unity coefficients. However, it yields a "very respectable" (considering its simplicity) low-pass frequency response which may be tuned by adjusting the filter length.

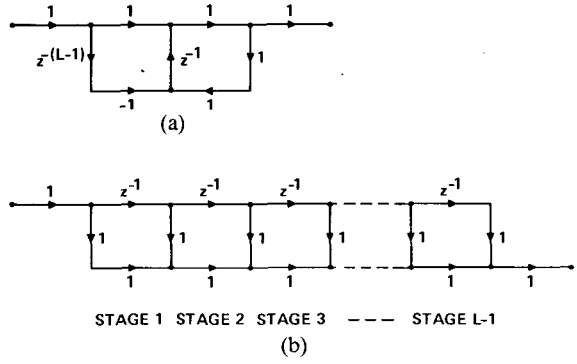


Fig. 1. (a) An efficient RRS prefilter realization. (b) Direct-form implementation of the RRS.

Due to its exceptionally low computational complexity, the RRS is an ideal candidate for use as a prefilter. As a specific example, we consider the following set of low-pass filter requirements: $\omega_p = 0.042\pi$ rad/sample, $\omega_s = 0.14\pi$ rad/sample, $DB_p = 0.2$ dB (passband ripple), $DB_s = 60$ dB (stopband attenuation). (ω_p and ω_s are the same as used for the examples given in [2].) By using the design rules presented in [1], it was determined that the optimum length for the RRS in this example is 14. The required length for the corresponding equalizer is 47. Fig. 2(a) shows the overall amplitude response of the prefilter–equalizer cascade. As a basis for comparison, a conventional filter was designed to meet the same specifications. The required

Manuscript received January 24, 1983. This work was supported by the National Science Foundation under Grant ECS82-06207.

J. W. Adams is with the Radar Systems Group of the Hughes Aircraft Company, El Segundo, CA 90009.
A. N. Willson, Jr. is with the Electrical Engineering Department, University of California, Los Angeles, CA 90024.

0098-4094/84/0300-0260$01.00 ©1984 IEEE

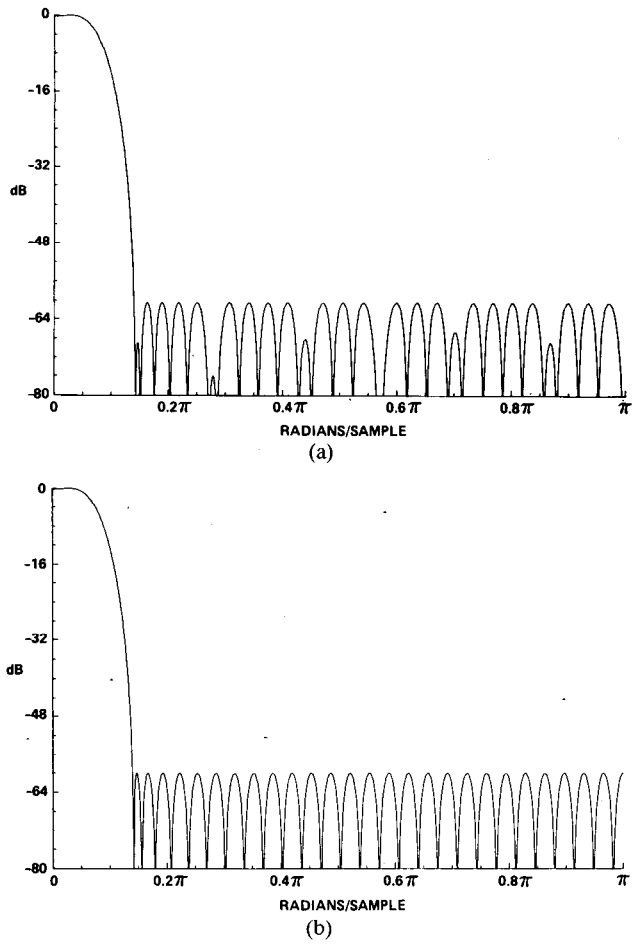


Fig. 2. (a) Amplitude response of the RRS-equalizer network. (b) Amplitude response of the conventional filter.

length was 59, substantially more than the length of the equalizer. Fig. 2(b) provides the amplitude response of the conventional filter. A summary of the hardware requirements for the RRS-equalizer cascade and the conventional filter is given in Table I. (Table I also contains the results for various other prefilter-equalizer designs which meet the same specifications. The other designs are discussed in subsequent sections.)

II. Cascaded RRS Prefilters

In Section I we discussed the use of a single RRS as a prefilter. Since the computational complexity of the RRS is very small, it would be worthwhile to consider using a prefilter composed of a cascade of RRS's. (Cascades of simple moving average filters were first exploited by Blackman and Tukey, and are discussed in [3].) Cascaded RRS's can provide increased stopband attenuation, which has the following benefits:

(1) Increased stopband attenuation in the prefilter effectively decreases the stopband attenuation requirements on the equalizer and allows its order to be reduced.
(2) Since the equalizer's coefficient quantization error spectrum is attenuated by the prefilter for frequencies in the stopband (as discussed in [1]), a reduced sensitivity to coefficient quantization can be expected.

TABLE I
Hardware Requirements Summary for the Example

	Delays	Adders	Multipliers	M + A
Conventional Filter	58	58	30	88
Single RRS Prefilter	14	2	0	
Equalizer	46	46	24	
Total	60	48	24	72
RRS Cascade Prefilter	24	4	0	
Equalizer	38	38	20	
Total	62	42	20	62
Modified Kaiser-Hamming	43	7	1	
Equalizer	30	30	16	
Total	73	37	17	54
Bateman-Liu Prefilter	17	8	0	
Equalizer	43	43	22	
Total	60	51	22	73

For the low-pass filter example discussed in Section I, we now consider using a second RRS in the prefilter. Clearly, the second RRS should be chosen on the basis of cancelling the sidelobe peaks of the first RRS. The peak of the highest (and first) sidelobe in an RRS with length 14 occurs at 0.6346 rad/sample. The first null in a length 10 RRS is at 0.6283 rad/sample. Fig. 3(a) provides overlays of the amplitude responses of RRS's with lengths of 14 and 10, while Fig. 3(b) shows the product of the individual responses. The length of the equalizer required for the RRS cascade is 39. The amplitude response of the prefilter–equalizer network is plotted in Fig. 3(c). The hardware requirements are given in Table I. Fig. 4 shows the amplitude responses of the conventional filter and the {RRS cascade}–equalizer network with coefficients rounded to 8 bits. As discussed in [1] the prefilter–equalizer network is much less sensitive to coefficient quantization.

A cascade of RRS's will have more stopband attenuation than a single RRS, but the passband rolloff will be increased. In the next section we consider a method which can improve the performance in both the stopband and the passband.

III. Prefilters Obtained from a Modification to the Kaiser–Hamming Filter Sharpening Method

If a filter is connected in series with a duplicate of itself, it is clear that the two-filter cascade will have twice as much (measured in decibels) stopband attenuation as each individual filter. But of course the passband ripple will be increased. Kaiser and Hamming [4] have devised a simple method for repeatedly using the same filter such that the overall frequency response is improved in both the passband and the stopband. They call this process "filter sharpening."

The basic structure for filter sharpening is shown in Fig. 5(a). Kaiser and Hamming refer to it as the SSS (simple

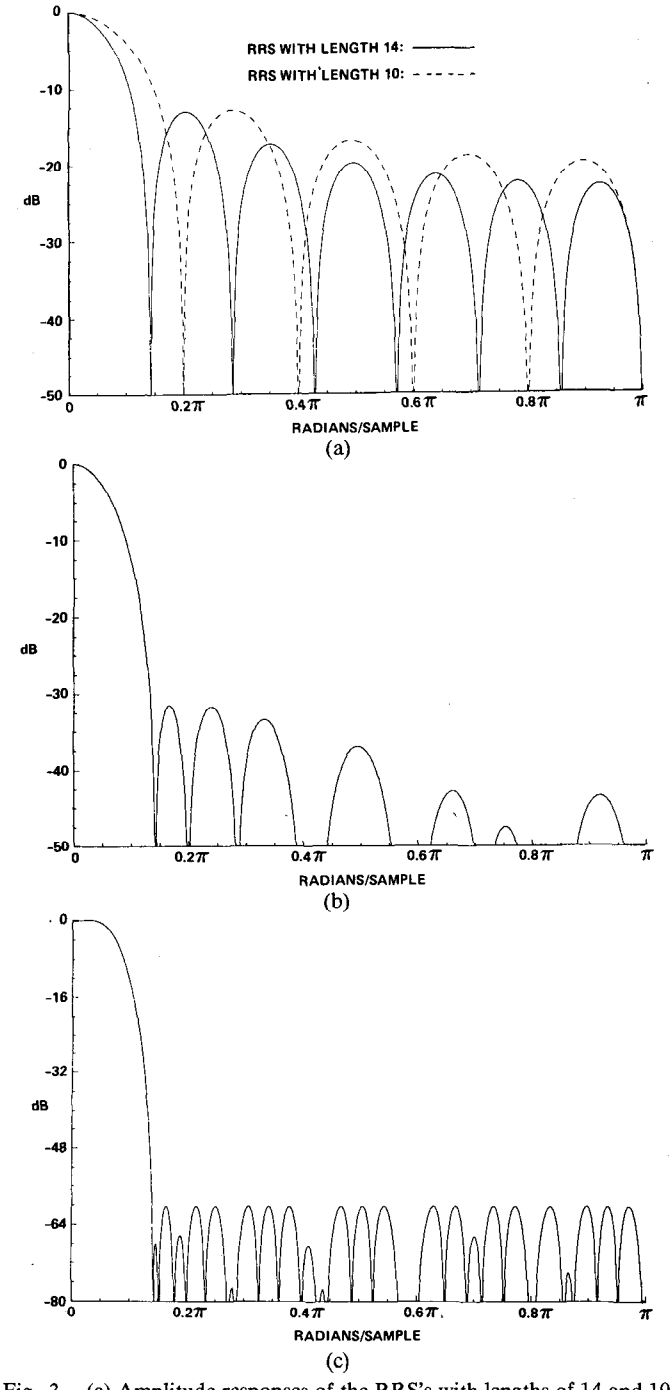


Fig. 3. (a) Amplitude responses of the RRS's with lengths of 14 and 10. (b) Amplitude response of the cascade of RRS's. (c) Amplitude response of the prefilter-equalizer network, where the prefilter is composed of the two RRS's in cascade.

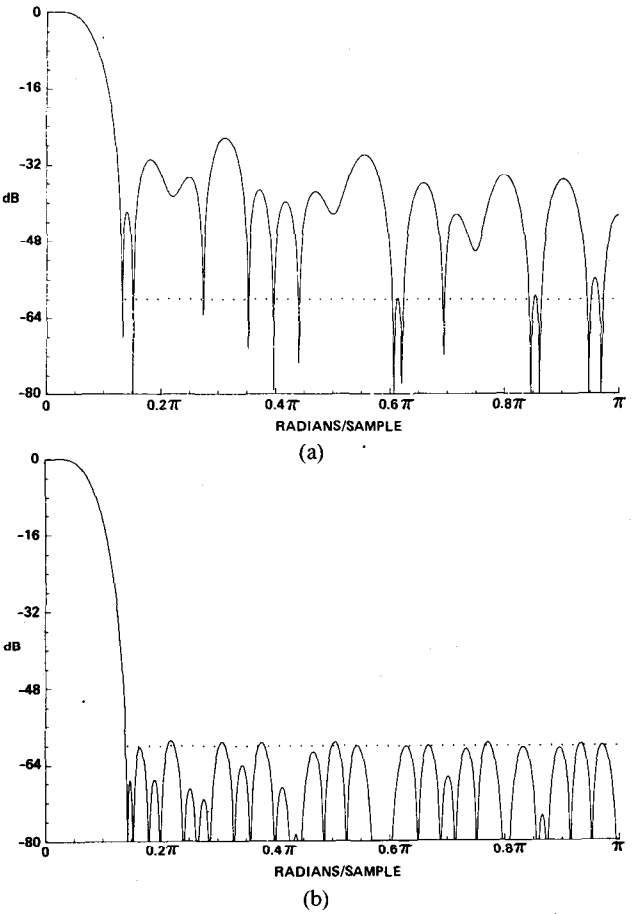


Fig. 4. (a) Amplitude response of the conventional filter with 8-bit coefficients. (b) Amplitude response of the {RRS cascade}-equalizer with 8-bit coefficients.

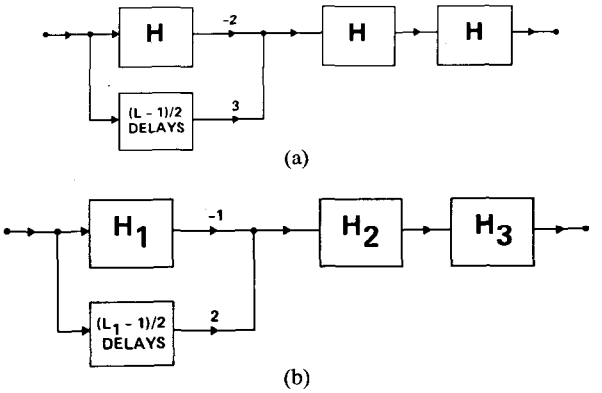


Fig. 5. (a) Kaiser–Hamming SSS structure. (b) Modified SSS structure.

symmetric sharpening) network. It uses the same filter (called the prototype) three times. The amplitude response of the SSS structure may be expressed as follows:

$$\bar{A}(e^{j\omega}) = A^2(e^{j\omega}) \cdot \{3 - 2A(e^{j\omega})\}$$

where

$A(e^{j\omega})$ = amplitude response of the prototype $= |H(e^{j\omega})|$.

In [4] Kaiser and Hamming derive the following approximations:

$$\bar{\delta}_p \simeq 3\delta_p^2 \qquad \text{(for small } \delta_p\text{)}$$

$$\bar{\delta}_s \simeq 3\delta_s^2 \qquad \text{(for small } \delta_s\text{)}$$

where

$\bar{\delta}_p, \bar{\delta}_s$ passband and stopband errors for the SSS network

δ_p, δ_s passband and stopband errors for the prototype filter.

For small values of δ_p and δ_s, the passband ripple and stopband attenuation of the SSS network are much better than for the prototype filter.

Modification of the SSS Network

In the stopband, the amplitude response of the SSS structure may be approximated by: $\bar{A}(e^{j\omega}) \simeq 3A^2(e^{j\omega})$. The SSS structure sacrifices stopband attenuation in two ways:

(1) Sidelobe peaks multiply sidelobe peaks and nulls multiply nulls for frequencies in the stopband. The structure would be more efficient if different filters were used such that the nulls of one cancelled the sidelobes of the other.

(2) The factor of 3 in $\bar{A}(e^{j\omega}) \simeq 3A^2(e^{j\omega})$ produces a 9.54-dB increase in the stopband gain over that of a simple two-filter cascade.

In order to alleviate these problems, the filter sharpening network shown in Fig. 5(b) is proposed. It has an amplitude response given by

$$\bar{A}(e^{j\omega}) = \{2 - A_1(e^{j\omega})\} A_2(e^{j\omega}) A_3(e^{j\omega}).$$

For frequencies in the stopband, the amplitude response may be approximated as: $\bar{A}(e^{j\omega}) \simeq 2A_2(e^{j\omega})A_3(e^{j\omega})$. Clearly, $A_2(e^{j\omega})$ and $A_3(e^{j\omega})$ should be chosen so that the nulls of one cancel the sidelobe peaks of the other. For the passband, the individual amplitude responses may be written as

$$A_i(e^{j\omega}) = 1 - \delta_i(e^{j\omega}), \qquad \text{for } i = 1,2,3$$

where $0 \leqslant \delta_i(e^{j\omega}) \ll 1$, and the peak gains of the prototype filters have been normalized to unity.

Then $\bar{A}(e^{j\omega})$ may be expressed as $\bar{A}(e^{j\omega}) = \{1 + \delta_1(e^{j\omega})\}\{1 - \delta_2(e^{j\omega})\}\{1 - \delta_3(e^{j\omega})\}$ for frequencies in the passband. Clearly, $\delta_1(e^{j\omega})$ should be chosen on the basis of compensating for the passband rolloff caused by $\delta_2(e^{j\omega})$ and $\delta_3(e^{j\omega})$.

Example

Suppose that a prefilter is needed for the example set of filter specifications discussed in Section I. An RRS with length 14 was used as the prefilter in Section I. Figs. 6 and 7 show the amplitude response of Kaiser and Hamming's original SSS network, using the length 14 RRS three times.

Now consider the sharpening structure which uses the proposed modification. Filters 2 and 3 should be chosen according to the same considerations which apply to the two-filter cascade discussed in Section II. Suppose RRS's with lengths of 14 and 10 are used, as for the example in Section II. The resulting rolloff at the passband edge due to filters 2 and 3 is -1.89 dB (a gain of 0.804).

In order to compensate for this we require $\{2 - A_1(e^{j\omega_p})\}\{0.804\} \simeq 1$. This implies that $A_1(e^{j\omega_p}) \simeq 0.756$. A good candidate for Filter 1 would be an RRS with length 19, which has a gain of 0.758 at ω_p. Figs. 6 and 7 show the amplitude response of our proposed sharpening network using RRS's with lengths of 19, 14, and 10. Clearly, the

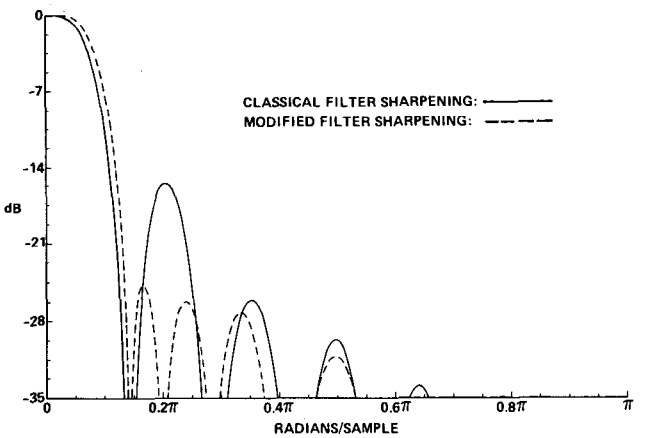


Fig. 6. *Solid Curve:* Amplitude response of the original SSS network using three RRS's, each with length 14. *Dashed Curve:* Amplitude response of the modified sharpening network using RRS's with lengths of 19, 14, and 10.

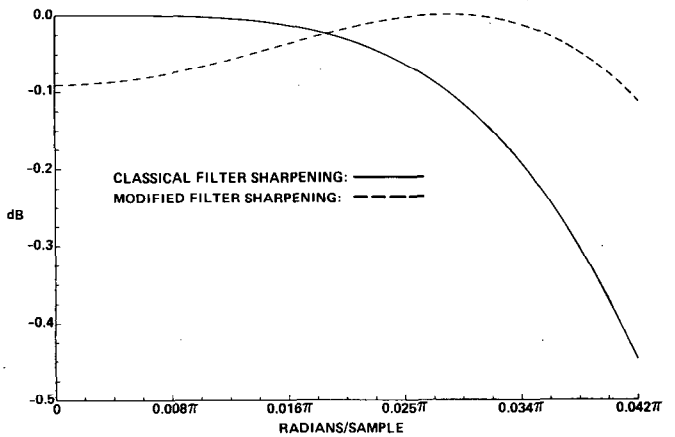


Fig. 7. Amplitude responses in the passband for the original SSS network and the modified filter sharpening network.

proposed structure provides a considerable improvement over the original SSS network. In the present context the "price" that one must pay for this improvement—i.e., the use of three different prototype filters rather than three copies of a single prototype—is inconsequential. (Notice that in terms of additional hardware requirements, only a single additional delay is needed, as $3 \times 14 = 42$, while $19 + 14 + 10 = 43$.)

An equalizer for the proposed sharpening network composed of RRS's with lengths of 19, 14, and 10 was designed to achieve the overall filter specifications previously mentioned. An equalizer of length 31 was required. The overall amplitude response is shown in Fig. 8. Table I summarizes the computational requirements, assuming that the $(L_1 - 1)/2$ delays shown in Fig. 5 are obtained by tapping into the first RRS, as shown in Fig. 9. It was also assumed for Table I that a multiplier would be involved in normalizing the dc gain of the first RRS to unity, since its length is not an integer power of 2. (It is important to normalize the peak gain of the first RRS to unity because of the parallel data path associated with it, as indicated in Fig. 5. The other two RRS's are merely in series, so that their scaling

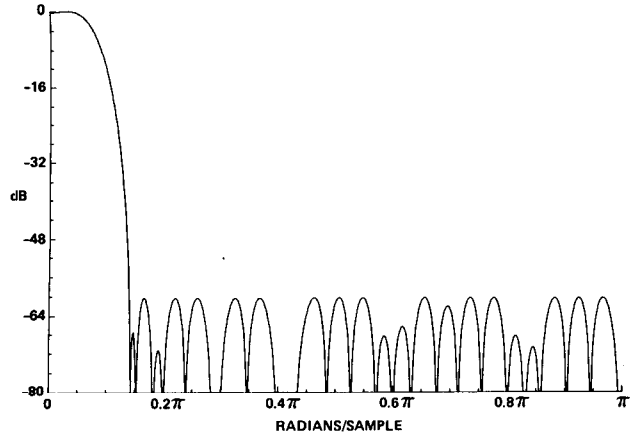


Fig. 8. Overall amplitude response of the modified sharpening network and its equalizer.

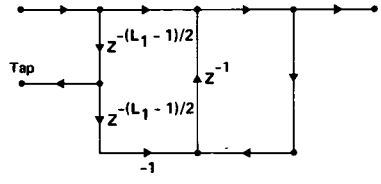


Fig. 9. A tap into the first RRS may be used to obtain the input signal delayed by $(L_1-1)/2$ samples.

does not affect the shape of the amplitude response of the SSS network. They may be scaled to a gain of approximately unity by using the appropriate number of bit shifts.)

IV. Prefilters Derived from the Method of Bateman and Liu

Bateman and Liu [2] have proposed an interesting approach to finite impulse response filter design for charge transfer devices (CTD's) which restricts all of the coefficients to have the values 0, +1, or −1, thus eliminating the need for multipliers. The Bateman–Liu filter has its origins in a communications technique called delta-modulation, as quoted from [2] in the following.

> "The basic idea behind our approach is to convert the filter coefficients into a sequence by a delta modulation-like scheme where instead of allowing only the values −1 and +1 we also allow the value 0. As in delta modulation, oversampling is required to keep the error at an acceptable level."

The same approach could also be used for *digital* filter design, except the following issues should be considered.

(1) It is difficult for the Bateman–Liu filter, using only $\{0, +1, -1\}$ as coefficients, to realize a frequency response having a small passband ripple and a large stopband attenuation. A very high filter order and a very large number of adders are needed to achieve even moderate levels of stopband attenuation.

(2) The sampling rate at which the Bateman–Liu filter operates typically must be much higher than the Nyquist rate. Digital interpolation and decimation may be used, but they can contribute to the computational complexity of the filter and also produce aliasing errors due to signals in the stopband folding back into the passband. Alternatively, if the A/D rate is increased, a more expensive A/D converter and faster signal processing hardware would be required.

TABLE II
Bateman–Liu Filter Design Examples ($\omega_p = 0.042\pi$ rad/sample, $\omega_s = 0.14\pi$ rad/sample)

	N	DB_p	DB_s	Resampling Requirements
Filter 1	17	1.4 dB	22 dB	None
Filter 2	102	0.5 dB	31 dB	Interpolation and decimation by a factor of 3.
Filter 3	292	0.2 dB	39 dB	Interpolation and decimation by a factor of 6.

In [2] Bateman and Liu provide three example filters with orders of 17, 102, and 292, all designed to have passband and stopband edge frequencies of 0.042π and 0.14π rad/sample. (For consistency we have used these same band edges for all of the examples in this paper.) Table II summarizes the performance of these Bateman–Liu filters. As can be seen from the table, extremely large increases in filter order are required to obtain relatively small improvements in passband ripple and stopband attenuation.

Instead of trying to meet a tight filter specification with the Bateman–Liu structure, we propose to avoid the above problems by relaxing the requirements on the Bateman–Liu filter and cascading it with an amplitude equalizer.

Example

Consider the filter specifications used for the example in Section I. ($\omega_p = 0.042\pi$ rad/sample, $\omega_s = 0.14\pi$ rad/sample, $DB_p = 0.2$ dB, $DB_s = 60$ dB). Filter 1 in the Bateman–Liu paper [2] has the same band edges, a reasonably low order (it uses only 8 adders), and does not require digital interpolation or decimation. We designed an equalizer to be used in conjunction with this filter; the required length was 44. The amplitude response of this {Bateman–Liu prefilter}–equalizer network is plotted in Fig. 10. The computational requirements are summarized in Table I.

Since the performance of the RRS was significantly improved when used in the modified Kaiser–Hamming SSS structure, it would be worthwhile to try using the Bateman–Liu filter in the same way. Unfortunately, it is very difficult to design Bateman–Liu filters for use in the modified filter sharpening network. The Bateman–Liu design procedure starts with a reference filter (typically computed by the Parks–McClellan computer program [5]) and then applies a delta-modulation process to it. The delta-modulation induces unpredictable errors which prevent the filter designer from conveniently synthesizing Bateman–Liu filters such that the nulls of one tend to cancel the sidelobe peaks of the other. In contrast, since a simple analytical

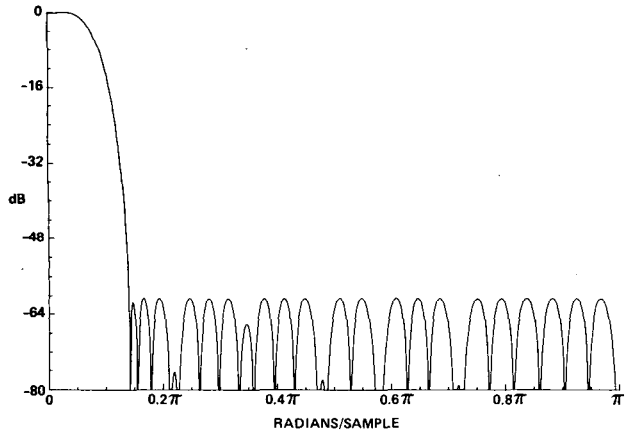


Fig. 10. Amplitude response of the {Bateman–Liu prefilter}–equalizer network.

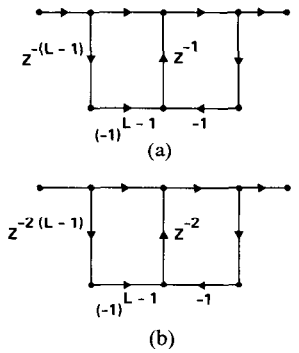


Fig. 11. (a) High-pass version of the RRS. (b) Bandpass version of the RRS with center frequency $\pi/2$.

expression exists for the frequency response of the RRS (as given in [1]), it is very easy to deal with in the context of the modified filter sharpening network.

V. High-Pass and Bandpass Filters

A low-pass filter may be converted to a high-pass filter by using the transformation: $z^{-1} \rightarrow -z^{-1}$. An efficient high-pass filter with passband and stopband edges of $\pi - \omega_p$ and $\pi - \omega_s$, respectively, may be obtained by applying the $z^{-1} \rightarrow -z^{-1}$ transformation to a low-pass prefilter–equalizer cascade with band edges of ω_p and ω_s. For the equalizer, the transformation $z^{-1} \rightarrow -z^{-1}$ is easily implemented by simply reversing the sign of every other coefficient. The RRS may be transformed as shown in Fig. 11(a).

An efficient bandpass filter with passband edges $\{\pi/2 - \omega_p/2, \pi/2 + \omega_p/2\}$ and stopband edges $\{\pi/2 - \omega_s/2, \pi/2 + \omega_s/2\}$ may be realized by applying the $z^{-1} \rightarrow -z^{-2}$ transformation to a low-pass prefilter–equalizer cascade with band edges of ω_p and ω_s. Fig. 11(b) shows the structure for the "bandpass RRS" with center frequency $\pi/2$ rad/sample. Fig. 12 gives a plot of the amplitude response of the example filter in Section III (a modified filter sharpening network using RRS's, followed by an equalizer) using the $z^{-1} \rightarrow -z^{-2}$ transformation. Bandpass versions of the RRS with center frequencies other than $\pi/2$ may also be obtained. They will be discussed in [6].

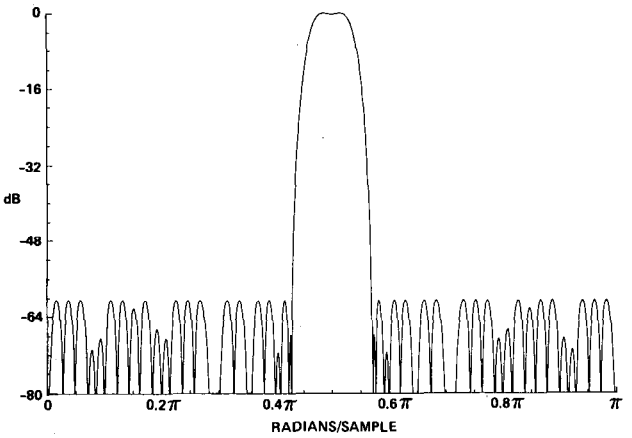


Fig. 12. Amplitude response of the example filter in Section III using the bandpass transformation: $z^{-1} \rightarrow -z^{-2}$.

References

[1] J. W. Adams and A. N. Willson, Jr., "A new approach to FIR digital filters with fewer multipliers and reduced sensitivity," *IEEE Trans. Circuits Syst.*, vol. CAS-30, pp. 277–283, May 1983.

[2] M. R. Bateman and B. Liu, "An approach to programmable CTD filters using coefficients 0, +1, and −1," *IEEE Trans. Circuits Syst.*, vol. CAS-27, pp. 451–456, June 1980.

[3] R. Blackman, *Data Smoothing and Prediction*. Reading, MA: Addison-Wesley, 1965.

[4] J. F. Kaiser and R. W. Hamming, "Sharpening the response of a symmetric non-recursive filter by multiple use of the same filter," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-25, pp. 415–422, Oct. 1977.

[5] J. H. McClellan, T. W. Parks, and L. R. Rabiner, "A computer program for designing optimum FIR linear phase digital filters," *IEEE Trans. Audio Electroacoust.*, vol. AU-21, pp. 506–526, Dec. 1973.

[6] J. W. Adams and A. N. Willson, Jr., "A novel approach to the design of efficient FIR digital bandpass filters," in *Proc. IEEE Int. Symp. on Circuits and Systems*, 1984.



John W. Adams (S'75–M'81) was born in Santa Monica, CA, on February 8, 1954. He received the B.S., M.S., Engr., and Ph.D. degrees in electrical engineering from the University of California, Los Angeles, in 1976, 1976, 1978, and 1982, respectively.

In 1978 he joined the Radar Systems Group of the Hughes Aircraft Company, where he now holds the position of Senior Staff Engineer in the Systems Engineering Department. His current research interests are in the areas of digital signal processing and synthetic aperture radar.

Dr. Adams is a member of Phi Beta Kappa, Tau Beta Pi, and Sigma Xi.



Alan N. Willson, Jr. (S'66–M'67–SM'73–F'78) was born in Baltimore, MD, on October 16, 1939. He received the B.E.E. degree from the Georgia Institute of Technology, Atlanta, GA, in 1961, and the M.S. and Ph.D. degrees from Syracuse University, Syracuse, NY, in 1965 and 1967, respectively.

From 1961 to 1964 he was with the IBM Corporation, in Poughkeepsie, NY. He was an Instructor in Electrical Engineering at Syracuse University from 1965 to 1967. From 1967 to 1973 he was a Member of the Technical Staff at Bell Laboratories, Murray Hill, NJ. Since 1973 he has been on the faculty of the University of California, Los Angeles, where he is now Professor of Engineering and

Applied Science, in the Electrical Engineering Department. In addition, he served the UCLA School of Engineering and Applied Science as Assistant Dean for Graduate Studies, from 1977 through 1981. He has been engaged in research concerning the stability of distributed circuits, properties of nonlinear networks, theory of active circuits, digital signal processing, and analog circuit fault diagnosis. He is editor of the book *Nonlinear Networks: Theory and Analysis*, IEEE Press, 1974.

Dr. Willson is a member of Eta Kappa Nu, Sigma Xi, Tau Beta Pi, the Society for Industrial and Applied Mathematics, and the American Society for Engineering Education. From June 1977 to June 1979 he served as Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS, and during 1980 he was Vice-President of the IEEE Circuits and Systems Society. He is currently serving as President of the IEEE Circuits and Systems Society. He is the recipient of the 1978 Guillemin–Cauer Award of the IEEE Circuits and Systems Society, for co-authoring the best paper published in their TRANSACTIONS during the previous year. He is the recipient of the 1982 George Westinghouse Award of the ASEE, and the 1982 Distinguished Faculty Award of the UCLA Engineering Alumni Association.

Maximum Amplitude Zero-Input Limit Cycles in Digital Filters

DAVID C. MUNSON, JR., MEMBER, IEEE, JAMES H. STRICKLAND, JR., MEMBER, IEEE, AND TIMOTHY P. WALKER, MEMBER, IEEE

Abstract—**An algorithm is developed for finding maximum amplitude zero-input limit cycles in recursive digital filters. Computer results are given which establish the feasibility of our algorithm and determine the tightness of previously published limit-cycle bounds. It is also shown that the maximum amplitude limit-cycle problem can be formulated as an integer linear program.**

I. INTRODUCTION

IT IS WELL KNOWN that with zero input, the output of a recursive digital filter may be nonzero due to arithmetic rounding or truncation. Since a digital filter is a finite state machine, such nonzero outputs must be periodic, and hence they have been termed zero-input limit cycles. Filter structures and rounding strategies have been developed for suppressing limit cycles, but due to hardware considerations, recursive filters are still often implemented as parallel or cascaded direct form second-order sections with standard rounding or truncation.

We consider direct form second-order filters with fixed-point arithmetic. For these filters, many limit-cycle properties have been discovered [1]–[5] and upper bounds on limit-cycle amplitudes have been derived [1], [6]–[10]. Unfortunately, though, the tightness of available limit-cycle bounds has not been firmly established and no algorithm has been suggested for finding exact maximum amplitude limit cycles. In cases where limit-cycle bounds indicate the *possibility* of large limit cycles, such an algorithm would help filter designers determine whether to use a low-error filter structure [11]–[14] or a special arithmetic [15], [16].

In this paper, we develop a practical algorithm for finding maximum amplitude limit cycles. This is accomplished by first developing an algorithm for finding maximum amplitude limit cycles of a particular period. This more restricted problem is solved by essentially considering each initial filter state that can be a maximum amplitude state of a limit cycle. A bound on the maximum period of a limit cycle is then derived so that this algorithm can be applied to determine the maximum amplitude limit cycle over all possible periods.

Although the resulting algorithm is quite straightforward, it is shown to be effective. Numerous computer results are tabulated which establish the feasibility of our approach and determine the tightness of the Long and Trick limit-cycle bounds [8]. Finally, we show that the maximum amplitude limit-cycle problem can be formulated as an integer linear program. This formulation is related to earlier work by Parker and Hess [6].

II. ALGORITHM FOR FINDING MAXIMUM AMPLITUDE LIMIT CYCLES

We consider the second-order filters with zero input in Figs. 1 and 2. Fig. 1 assumes quantization after the accumulator and Fig. 2 shows quantization after each multipli-

Manuscript received February 15, 1982; revised August 27, 1982. This work was supported in part by the Joint Services Electronics Program under Contract N00014-79-C-0424. Portions of this paper were presented at the 1981 IEEE International Conference on Acoustics, Speech, and Signal Processing (Atlanta, GA), March 30–April 1, 1981.

D. C. Munson, Jr. is with the Coordinated Science Laboratory and the Department of Electrical Engineering, University of Illinois, Urbana, IL 61801.

J. H. Strickland, Jr. was with Bell Laboratories, Holmdel, NJ 07733. He is now at 3808 Montevallo Road, Birmingham, AL 35213.

T. P. Walker is with Teradyne, Inc., Boston, MA 02111.

0098-4094/84/0300-0266$01.00 ©1984 IEEE