The IEEE International Midwest Symposium on Circuits and Systems (MWSCAS 2002)
Tulsa OK, USA, August 4–7, 2002.

A DOUBLE BINARY-TREE FIR-FILTER FORM AND THE CORRESPONDING LOW-SENSITIVITY IIR BANDPASS/BANDSTOP STRUCTURE

Jeffrey O. Coleman

Naval Research Laboratory
Washington, DC
http://alum.mit.edu/www/jeffc

***Abstract*— An FIR filter comprising back-to-back binary trees, each using multiple-sample delay blocks, offers advantages (1) for rate-changing filters, (2) when block delays are especially hardware efficient, and (3) when replacing block delays with allpass sections to create an IIR filter. In the latter case, the peak group delays of the allpass sections approximate the several distinct delays of the blocks replaced. With allpass sections fixed, the combining weights are easily optimized.**

1 INTRODUCTION

Hardware implementation of an N-th order FIR filter traditionally involves N unit delays z^{-1}, with the output of each used in power-consuming computation on every clock cycle. Using multi-sample delay blocks of form z^{-m} instead would yield power savings if the total number of blocks were significantly less than N, since only block outputs would enter into computations or otherwise effect state changes on a given clock cycle. Circular addressing would presumably keep stationary all data internal to the block. This paper first presents an FIR-filter topology based on just such a use of multi-sample block delays.

The paper then briefly considers replacing block delays with nonidentical allpass sections, with the "FIR" coefficients then optimized to shape the overall frequency response. The resulting IIR filter can be of extremely high order without introducing stability or coefficient-sensitivity problems.

2 FIR-FILTER FORMS

Direct, Transposed, and Polyphase Forms

A quick review of standard FIR-filter structures will establish the notation to be used later to introduce new ones.

Use an $L \times L$ identity matrix $\mathbf{I}$ to define matrix $\mathbf{z}_m$ as

$$\mathbf{z}_m \triangleq \begin{pmatrix} \mathbf{I} & & & \\ & z^{-m}\mathbf{I} & & \\ & & z^{-2m}\mathbf{I} & \\ & & & \vdots \\ & & & z^{-(N-1)m}\mathbf{I} \end{pmatrix}, \qquad (1)$$

This work was supported by the Office of Naval Research.

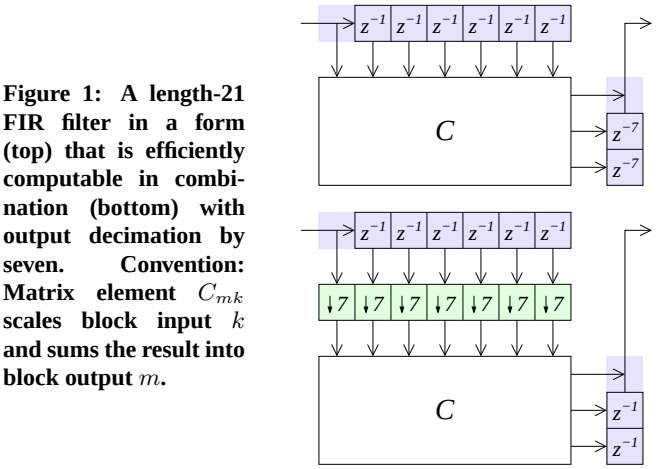

Figure 1: A length-21 FIR filter in a form (top) that is efficiently computable in combination (bottom) with output decimation by seven. Convention: Matrix element C_{mk} scales block input k and sums the result into block output m.

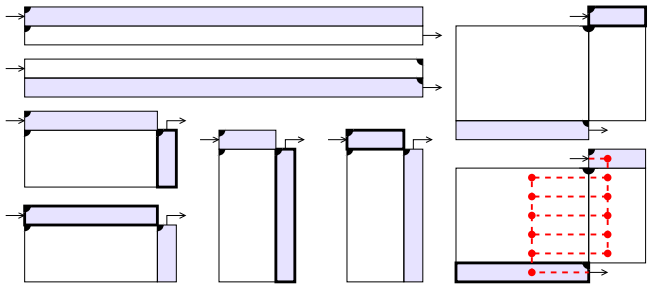

Figure 2: FIR filter structures in a streamlined Fig. 1 notation. One delay vector (nonbold) in each structure is M-vector $\mathbf{z}_1$ and the other (bold) is $\mathbf{z}_M$. Dots mark the zero-index corners of coefficient blocks and the zero-delay ends of delay blocks. Using two coefficient blocks (right) yields a dot-product (dashed lines) net coefficient at each net delay.

and for now let $L = 1$ so that $\mathbf{I} = 1$ with $\mathbf{z}_m$ an N-vector. Then in *direct form* an FIR filter realizes $Y(z) = \mathbf{c}^T\mathbf{z}_1 X(z)$ from coefficient vector $\mathbf{c}$ in a right-associative way, and in *transposed form* it instead realizes $Y(z) = \mathbf{z}_1^T\mathbf{c}X(z)$.

Next suppose coefficient matrix $\mathbf{C}$ has M columns. Then $Y(z) = \mathbf{z}_M^T\,\mathbf{C}\,\mathbf{z}_1 X(z)$ realizes a filter with length a multiple of M as illustrated in Fig. 1 for $M = 7$. Output decimation by M need not waste computation, as it can be referred to the coefficient-block input to be combined with input delay vector $\mathbf{z}_1$ to become a polyphase decomposition. Output delay

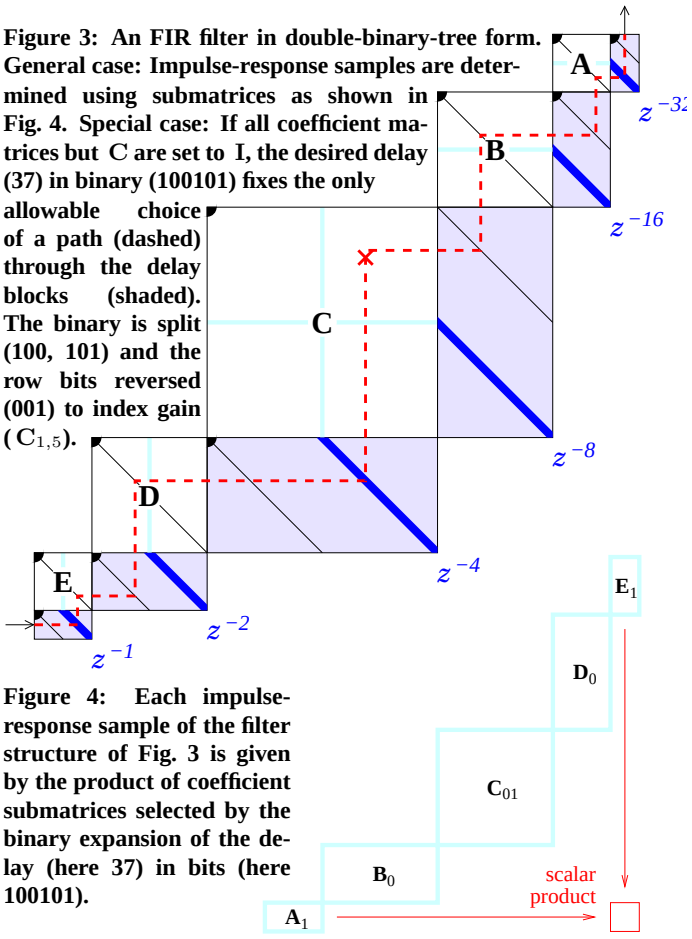

Figure 3: An FIR filter in double-binary-tree form. General case: Impulse-response samples are determined using submatrices as shown in Fig. 4. Special case: If all coefficient matrices but $\mathbf{C}$ are set to $\mathbf{I}$, the desired delay (37) in binary (100101) fixes the only allowable choice of a path (dashed) through the delay blocks (shaded). The binary is split (100, 101) and the row bits reversed (001) to index gain ($\mathbf{C}_{1,5}$).

Figure 4: Each impulse-response sample of the filter structure of Fig. 3 is given by the product of coefficient submatrices selected by the binary expansion of the delay (here 37) in bits (here 100101).

$\mathbf{z}_M$ becomes $\mathbf{z}_1$, and the coefficient block and output delay together become an FIR filter in *transposed polyphase form*.

Similarly, structure $Y(z) = \mathbf{z}_1^T\,\mathbf{C}\;\mathbf{z}_M X(z)$ with M-row coefficient matrix $\mathbf{C}$ allows input interpolation to be referred to the matrix output, turning input delay $\mathbf{z}_M$ into $\mathbf{z}_1$. The interpolator and output delay together realize polyphase reconstruction. The rest is an FIR filter in *direct polyphase form*.

Figure 2 shows the filter structures discussed thus far and new ones that compute $Y(z) = \mathbf{z}_1^T\,\mathbf{C}_2\,\mathbf{C}_1\mathbf{z}_M X(z)$ (top) and $Y(z) = \mathbf{z}_M^T\,\mathbf{C}_2\,\mathbf{C}_1\mathbf{z}_1 X(z)$ (bottom). Here a two-matrix product replaces the single coefficient matrix used earlier.

The Double Binary-Tree Form

The two rightmost structures in Fig. 2 transform a scalar signal into a vector in the input-delay stage, use matrix multiplication to increase the vector's length twice, and finally reduce the vector back to a scalar in the output-delay stage. The structure in Fig. 3 likewise manages signal-vector length with matrix products but does so more systematically. It increases signal-vector length from unity to some maximum using a binary tree of length-doubling delay stages and then reduces it back to unity output using a reversed binary tree of length-halving delay stages. Altogether it realizes

$$Y(z) = \mathbf{z}_{32}^T\,\mathbf{A}\,\mathbf{z}_{16}^T\,\mathbf{B}\,\mathbf{z}_8^T\,\mathbf{C}\,\mathbf{z}_4\,\mathbf{D}\,\mathbf{z}_2\,\mathbf{E}\,\mathbf{z}_1 X(z). \qquad (2)$$

stage	*elements*	*matrix*	*delays*	*matrix*	*delays*
input	1	$\mathbf{z}_1$	1	$\mathbf{z}_{32}$	32
2nd	2	$\mathbf{z}_2$	4	$\mathbf{z}_8$	16
3rd	4	$\mathbf{z}_4$	16	$\mathbf{z}_2$	8
4th	4	$\mathbf{z}_8$	32	$\mathbf{z}_1$	4
5th	2	$\mathbf{z}_{16}$	32	$\mathbf{z}_4$	8
output	1	$\mathbf{z}_{32}$	32	$\mathbf{z}_{16}$	16
		total:	117	*total:*	84

Table 1: In Fig. 3 each delay *stage* has some n_e *elements* z^{-m} in its *matrix* and so requires $n_e m$ *delays*. The Fig. 3 ordering of stages (center) is not optimal. One alternative is shown (right).

Here size-change factor $N = 2$ in (1) but L is arbitrary, so

$$\mathbf{z}_m = \begin{pmatrix} \mathbf{I} \\ \mathbf{I}\,z^{-m} \end{pmatrix} \qquad (3)$$

and $(\mathbf{z}_m)_{Ln+\ell,k} = z^{-mn}\,\mathbf{I}_{\ell,k} = z^{-mn}\delta_{\ell,k}$ using zero-origin array indexing with $n \in \{0,1\}$ and $0 \le \ell < L$. Splitting into submatrices in (2) yields three partial-product types:

$$\mathbf{z}_{16}^T\,\mathbf{B} = \begin{pmatrix} \mathbf{I} & z^{-16}\,\mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{B}_0 \\ \mathbf{B}_1 \end{pmatrix} = \mathbf{B}_0 + \mathbf{B}_1\,z^{-16}$$

$$\mathbf{D}\;\mathbf{z}_2 = \begin{pmatrix} \mathbf{D}_0 & \mathbf{D}_1 \end{pmatrix} \begin{pmatrix} \mathbf{I} \\ \mathbf{I}\,z^{-2} \end{pmatrix} = \mathbf{D}_0 + \mathbf{D}_1\,z^{-2}$$

$$\begin{aligned} \mathbf{z}_8^T\,\mathbf{C}\;\mathbf{z}_4 &= \begin{pmatrix} \mathbf{I} & z^{-8}\,\mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{C}_{00} & \mathbf{C}_{01} \\ \mathbf{C}_{10} & \mathbf{C}_{11} \end{pmatrix} \begin{pmatrix} \mathbf{I} \\ \mathbf{I}\,z^{-4} \end{pmatrix} \\ &= \mathbf{C}_{00} + \mathbf{C}_{01}\,z^{-4} + \mathbf{C}_{10}\,z^{-8} + \mathbf{C}_{11}\,z^{-12}. \end{aligned}$$

Partial products of this form then combine in a general result:

$$\frac{Y(z)}{X(z)} = \sum_{r,s,t,u,v,w \in \{0,1\}} \mathbf{A}_r\,\mathbf{B}_s\,\mathbf{C}_{tu}\,\mathbf{D}_v\,\mathbf{E}_w\,z^{-(32r+16s+8t+4u+2v+w)} \qquad (4)$$

A scalar impulse-response sample is the product of submatrices selected by the sample's delay in binary.

As a simple special case, let $\mathbf{A}$, $\mathbf{B}$, $\mathbf{D}$, and $\mathbf{E}$ be identity matrices and use (3) on 8×4 matrix $\mathbf{z}_8$ in (2) so that

$$\begin{aligned} \left(\mathbf{z}_{32}^T\,\mathbf{z}_{16}^T\,\mathbf{z}_8^T\right)_{4t+2s+r} &= \sum_{k=0}^{3} (\mathbf{z}_8)_{4t+2s+r,\,k}\,(\mathbf{z}_{16}\,\mathbf{z}_{32})_k \\ &= z^{-8t}\,(\mathbf{z}_{16}\,\mathbf{z}_{32})_{2s+r}\,. \end{aligned}$$

Continuing, and subscripting $\mathbf{C}$ to select elements now,

$$\frac{Y(z)}{X(z)} = \sum_{r,s,t,u,v,w \in \{0,1\}} \mathbf{C}_{4t+2s+r,\,4u+2v+w}\,z^{-(32r+16s+8t+4u+2v+w)}$$

Bit-reversal function $\rho(4r+2s+t) \triangleq 4t+2s+r$ for $r,s,t \in \{0,1\}$ leads to a convenient, more compact form:

$$\frac{Y(z)}{X(z)} = \sum_{u,v \in \{0,\ldots,7\}} \mathbf{C}_{\rho(u),\,v}\,z^{-(8u+v)}. \qquad (5)$$

In configurations in Fig. 2 that split the net delay, either delay block can be labeled with $\mathbf{z}_M$ and the other labeled $\mathbf{z}_1$.

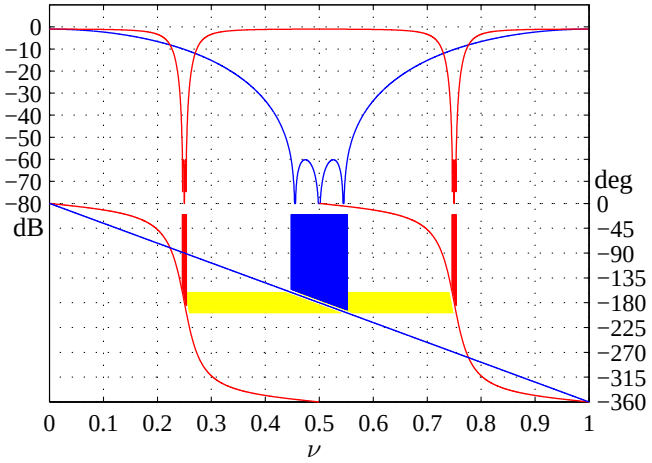

Figure 5: Frequency warping. A four-tap linear-phase lowpass FIR prototype filter (magnitude response above) is transformed into a bandstop IIR filter (magnitude response above) by replacing the prototype's unit delays (phase response below) with identical second-order allpass sections (phase response below). A common phase range (deg. axis below) relates the stopbands (freq. ν axis below). Allpass poles are as shown in Fig. A3.

These two labeling options become many in Fig. 3. Labels $\{\mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_4, \mathbf{z}_8, \mathbf{z}_{16}, \mathbf{z}_{32}\}$ can be assigned to stages in 6 ! ways, and the total number of delay units is always identical here for $2^{6/2}$ or eight of those ways. Table 1 compares one of the eight optimal assignments to the one shown in Fig. 3.

Separate realization of z^{-1} factors is wasteful. The Virtex FPGA [1] design library includes an efficient multi-sample delay z^{-m} for $m \leq 16$ based on circular table addressing. In VLSI, multi-sample delays could lower the memory-access rate and hence power consumption.[1] These ideas are not explored further here. In the sequel z^{-m} is instead approximated with an allpass section to create an IIR filter.

Transfer function (2) and its derivation show that the impulse response of a double binary-tree FIR filter is of some length 2^{2M} and characterized by $2M-1$ square coefficient matrices. (Less-symmetric tree topologies can realize other impulse-response lengths.) The central coefficient matrix, here $\mathbf{C}$, is $2^M \times 2^M$. For now the (5) special case is assumed.

3 ALLPASS SECTIONS FOR DELAY

The conventional replacement of all z^{-1} delays with identical IIR allpass sections is reviewed next. The central idea of this paper is then examined: Use nonidentical allpass sections instead, choosing allpass parameters according to the delay m of the multi-sample delay z^{-m} replaced.

Frequency-Warping Transformation

Building a tapped-delay-line filter from allpass delay units having frequency response $e^{j\phi(\nu)}$ results in frequency response

[1] Thanks to Oscar Gustafsson at Linköping University for pointing this out.

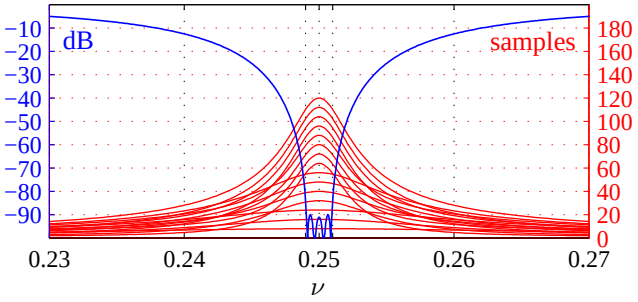

Figure 6: Group delays of 16 paths through allpass sections with $\sigma = 0$, $\gamma \in \{1, \frac{1}{2}, \frac{1}{4}, \frac{1}{8}\}$. IIR bandstop-filter weights minimize MS |gain| in a stopband bounded by $\nu = 0.25 \pm 0.001$ while fixing to unity MS |gain| in a reference band $\nu \in [0.001, 0.1]$.

form $\sum_n w_n e^{j\phi(\nu)n}$, the response $\sum_n w_n e^{-j2\pi\nu n}$ of a "prototype" FIR filter but with radian frequency $-2\pi\nu$ replaced by $\phi(\nu)$ to effect the frequency-axis mapping

$$\nu \rightarrow \{\eta\colon \phi(\eta) = -2\pi\nu\} = \phi^{-1}(\{-2\pi\nu\}).$$

Figure 5 illustrates this standard textbook idea. There the "mapped-to" set $\phi^{-1}(\{-2\pi\nu\})$ is periodic with unit period and is symmetric about each of the frequencies 0 and 0.5. Applying $d/d\nu$ to $\phi(\eta) = -2\pi\nu$ and solving yields

$$\frac{d\nu}{d\eta} = -\frac{1}{2\pi}\frac{d\phi(\eta)}{d\eta},$$

so the frequency-axis compression factor is just the group-delay response of the allpass section. In the examples to follow the second-order allpass section of Appendix A is used exclusively, but none of the ideas is specific to that choice.

A Double Binary Tree of Nonidentical Allpass Sections

If each z^{-m} block delay is replaced with an allpass section of peak group delay $\propto m$, does frequency warping predict the result? Near the pole frequency, yes. Elsewhere, no.

An experiment: Use a 4×4 coefficient matrix in the center and identity matrices elsewhere, so that six allpass sections of four distinct designs are required. The poles of the second-order section in Appendix A have radii $\approx 1-\gamma/2$ for $\gamma \ll 1$ and pole frequencies ≈ 0.25, so the computationally efficient form $\gamma = 2^{-M}$ yields the desired group-delay ratios. Group delay curves for the 16 paths associated with coefficient-matrix elements appear in Fig. 6. Because curve-spacing uniformity deteriorates rapidly as ν is moved from 0.25, it is far from clear that a coefficient matrix resulting in a desirable filter response over the whole frequency axis is possible. An optimization experiment is called for.

The bandstop filter of Fig. 6 was designed using a simple generalized-eigenvalue method [2] to optimize the 16 elements of the coefficient matrix. (Second-order cone programming, the method of weighted least squares, etc., could just as well have been used.) Clearly, this single, preliminary experiment yields a reasonable solution, but just as clearly, whether this approach is broadly useful remains to be seen.

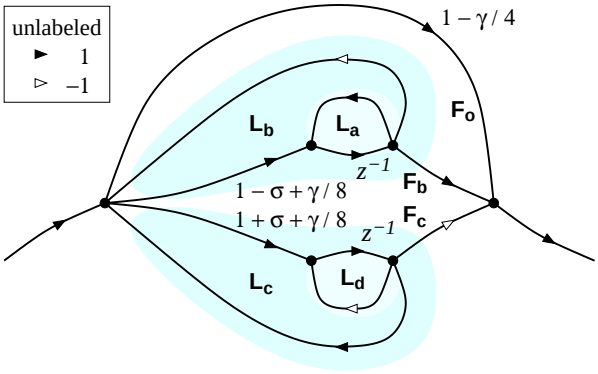

Figure A1: The transfer function of the signal-flow graph above is easily derived from the gains of its (shaded) loops L_a, L_b, L_c, and L_d, and its forward paths F_o, F_b, and F_c.

4 IN LIEU OF A CONCLUSION: AN OPEN PROBLEM

Might the FIR coefficients be systematically factored into the matrix-product form of (4) using matrices with simple elements, like zeros or powers of two? The potential for extraordinary computational efficiency is quite appealing.

Appendix A: A SECOND-ORDER ALLPASS SECTION

Summerfield [3] presented a second-order allpass section that was based on a three-port adapter and designed to minimize the critical-path delay. However, in some applications the entire allpass section is itself inside an enclosing loop. The maximum-length unregistered forward path, four additions in Summerfield's section, then becomes part of the critical loop of the larger system. The modified form of Summerfield's section derived here addresses such applications by minimizing the maximum unregistered forward-path delay.

Begin with the signal-flow graph of Fig. A1, the transfer function of which is given by

$$H(z) = \frac{z^{-2} - 2\sigma z^{-1} + \quad 1 - \gamma/4}{1 - 2\sigma z^{-1} + (1-\gamma/4)z^{-2}} \ , \qquad \text{(A1)}$$

an allpass response with pole locations specified by the peculiar coordinates—a reason for this will appear in due course—of Fig. A2. This places poles at $\sigma \pm j\sqrt{1-(\gamma/4)-\sigma^2}$, and the characteristic frequencies of the frequency-warping transformation realized by this section are then just the pole angles $\{\omega_c\}$ defined by $\cos\omega_c = \sigma/\sqrt{1-\gamma/4}$.

Why the (σ, γ) parameterization? Most applications require the frequency-warping curve in Fig. 5 to be steep, so the poles must lie near the unit circle. Often the pole radius need not be set precisely, however, as within reasonable limits the prototype filter can be designed in accordance with whatever warping curve is convenient. So for many applications γ can

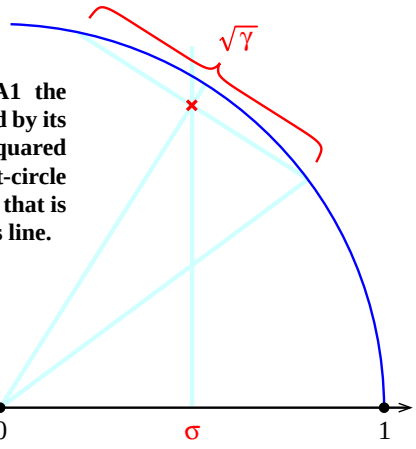

Figure A2: In Fig. A1 the pole location is specified by its real part σ and the squared length γ of the unit-circle chord through the pole that is orthogonal to its radius line.

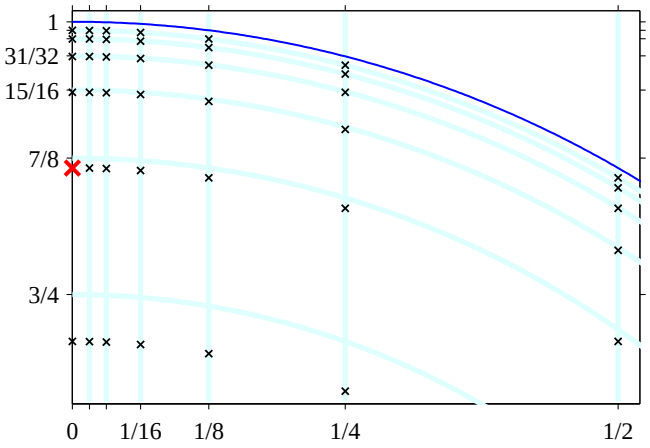

Figure A3: The Fig. A1 allpass section's pole locations when σ and γ of Fig. A2 are set to various powers of two (including $\sigma = 2^{-\infty}$). The $\sigma = 0$, $\gamma = 1$ pole used in Fig. 5 is emphasized.

be set to a (possibly negative) power of two. In cases in which γ is thereby made small, the poles are separated from the unit circle by distance $1 - \sqrt{1-\gamma/4} \approx \gamma/2$, yet another power of two. Finally, σ can be set to the sum of some small number of power-of-two terms, possibly as few as one term or even none at all—this is the $\omega_c = \pm\pi/2$ case—as for the possible pole locations illustrated in Fig. A3.

There are many other low-delay section designs, *e.g.* [4].

References

[1] Xilinx, Inc. http://support.xilinx.com/partinfo/databook.htm.

[2] J. O. Coleman, "Quadratic FIR-filter design as a generalized eigenproblem," in *1998 Eureopean Signal Processing Conf., Rodos, Greece*.

[3] M. S. Anderson, S. Summerfield, and S. S. Lawson, "Realisation of lattice wave digital filters using three-port adaptors," *Electronics Letters*, vol. 31, pp. 628–629, 13 April 1995.

[4] M. Bhattacharya and T. Saramäki, "Multiplierless implementation of bandpass and bandstop recursive digital filters," in *2002 IEEE Int'l Symp. on Circuits and Systems, Phoenix AZ, USA.*