MAGNITUDE SQUARED DESIGN OF RECURSIVE FILTERS WITH THE CHEBYSHEV NORM USING A CONSTRAINED RATIONAL REMEZ ALGORITHM

Ivan W. Selesnick, Markus Lang and C. Sidney Burrus
Department of Electrical and Computer Engineering,
Rice University, Houston, TX 77251-1892

May 1994

Abstract

We describe a Remez type exchange algorithm for the design of stable recursive filters for which the Chebyshev norm of $H(\omega) - F(\omega)$ is minimized, where $H(\omega)$ and $F(\omega)$ are the realized and desired magnitude squared frequency responses. The number of poles and zeros can be chosen arbitrarily and the zeros do not have to lie on the unit circle. The algorithm allows us to design filters with non-conventional frequency responses with arbitrary weighting functions. It also gives optimal minimum phase FIR filters and Elliptic recursive filters as special cases. We discuss three main difficulties in the use of the Remez algorithm for recursive filter design and give ways to overcome them.

EDICS: SP 2.1.2 General IIR Systems

Contact Address:
Ivan W. Selesnick
Electrical and Computer Engineering - MS 366
Rice University,
Houston, TX 77251-1892
Phone: (713) 527-8101 x3508
FAX: (713) 524-5237
email: selesi@jazz.rice.edu

`Submitted to:  IEEE Transactions on Signal Processing - May 1994`

List of Figures

List of Tables

1 Introduction

In this paper we describe a Remez exchange algorithm for the design of stable recursive digital filters in the frequency domain that does not require the phase of the desired frequency response. The number of poles and zeros can be chosen arbitrarily and the zeros do not have to lie on the unit circle. The design method constructs a filter whose magnitude squared frequency response is a Chebyshev approximation to a desired nonnegative function. In this way, optimal minimum phase FIR filters and Elliptic recursive filters can be designed as special cases.

The approximation algorithm we use minimizes the Chebyshev norm of $H(\omega) - F(\omega)$ where $H(\omega)$ and $F(\omega)$ are the realized and desired magnitude squared frequency responses respectively. Our approach constrains the approximation $H(\omega)$ to be nonnegative, for then it can be spectrally factored to obtain a stable filter whose magnitude squared frequency response approximates $F(\omega)$. To obtain these nonnegative Chebyshev approximations in the frequency domain we modify the rational Remez exchange algorithm described by Powell [**?**] (also see [**?**]).

It appears that the rational Remez exchange algorithm is infrequently used for the design of recursive digital filters. Among the possible reasons for this is the need to solve a set of nonlinear equations at each iteration. These equations have multiple solutions and are generally solved with Newton's method which may give a useless solution. However, by turning the nonlinear equations into a generalized eigenvalue problem one obtains every solution, only one of which (if any) is the appropriate one to use. Moreover, the appropriate solution can be chosen without ambiguity as explained below.

Another reason that the rational Remez exchange algorithm is not commonly used is the necessity that the magnitude squared approximation be nonnegative. Only then can the approximation be spectrally factored to obtain a stable filter. The relevant constrained approximation problem can, however, be solved as easily as the unconstrained problem. In fact in [**?**] it was shown how to incorporate upper and lower bound constraints in the Parks-McClellan FIR filter design program. Here, we describe the necessary modifications to the rational Remez algorithm so that best *nonnegative* rational approximations can be found.

Possibly the most important reason the rational Remez exchange algorithm has not been more widely used is that it is not guaranteed to converge. It fails to converge when all the solutions to the nonlinear equations associated with the interpolation step have denominators that are not strictly positive. This break down of the algorithm can occur under two situations, (i) the best approximation over the interval is degenerate (a pole-zero pair can be said to cancel), and (ii)

high sensitivity to the initial reference set. We suggest a method for overcoming this situation by perturbing the reference set appropriately. We have observed that it is sometimes necessary to change the reference only slightly to make the rational Remez converge successfully.

Some papers on the design of recursive filters according to the Chebyshev norm require all the zeros of the filter to lie on the unit circle [?, ?, ?, ?] or require a special form for the frequency response [?]. While the best Chebyshev approximation may indeed have all its zeros on the unit circle, they may not and the above referenced methods will then give sub-optimal solutions. For example, an optimal wide band low pass filter with only 2 poles will possess zeros off the unit circle. Deczky [?, ?] pose a general optimization procedure based on second order sections and hence these algorithm may converge to a local optimum. In [?] an algorithm is given which relies on the desired frequency response being bandpass so that the numerator and denominator can be treated independently. The differential correction algorithm used in [?] is a robust algorithm but is computationally intensive since it requires the solution to a sequence of linear programming problems and does not take advantage of the alternation property (see below). An earlier paper that also uses linear programming methods is [?].

2 The Rational Remez Exchange Algorithm

The Remez exchange algorithm for (unconstrained) Chebyshev approximation by rational functions is based on the alternation property and an interpolation step, as is the polynomial Remez algorithm. We use the notation,

$$H(\omega) = \frac{a_0 + a_1\cos(\omega) + \ldots + a_m\cos(m\omega)}{1 + b_1\cos(\omega) + \ldots + b_n\cos(n\omega)} \tag{1}$$

for the realized magnitude squared frequency response and denote the numerator and denominator by $A(\omega)$ and $B(\omega)$ respectively. Note that the number of degrees of freedom among the cosine coefficients is $m+n+1$. We call the set of all such functions $R_{m,n}$. We let $\overline{R}_{m,n}$ be the subset of $R_{m,n}$ for which the denominator has no zeros in $[0,\pi]$.

Let $S \subset [0,\pi]$ be a union of intervals and let $F(\omega)$ be the desired non-negative function. By the best rational Chebyshev approximation from $\overline{R}_{m,n}$ to $F(\omega)$ over S we mean the function $H(\omega)$ in $\overline{R}_{m,n}$ that minimizes

$$||H(\omega) - F(\omega)|| = \max_{\omega\in S} |H(\omega) - F(\omega)|.$$

For any approximation $H(\omega)$, we denote the error function $H(\omega) - F(\omega)$ by $E(\omega)$. Note that the denominator of the rational function that best approximates $F(\omega)$ in the Chebyshev norm must have no zeros in S. If it did, $H(\omega)$ would be infinite at these points and clearly it could not be the best Chebyshev approximation. But when S is composed of different intervals, as is often the case in the design of digital filters where S is made up of the frequency response bands, it is possible that the best approximation by a rational function has a pole in a transition band. To prevent this, we explicitly limit the set of rational function to those whose denominators are strictly positive on $[0, \pi]$. The inclusion of a weighting function for the error is straightforward and in example 3 below, a non-constant weighting function is used to illustrate this.

Alternation Property: Recall that best Chebyshev approximations by polynomials ($n = 0$) are uniquely characterized by an alternation property. That is, there must be $m + 2$ points where the absolute value of the error function attains its maximum value, and the sign of the error at these points must alternate. However, in the rational case, this condition is only sufficient [**?**]:

Theorem 1 *Let $(\omega_1, ..., \omega_{n+m+2})$ be a sequence of points of S in ascending order (a reference set), and let $F(\omega)$ be a continuous function on S. If $H(\omega)$ is in $\overline{R}_{m,n}$ and if the equations*

$$H(\omega_i) + (-1)^i \delta = F(\omega_i) \tag{2}$$

for $i = 1, ..., m+n+2$ hold for $|\delta| = ||H(\omega) - F(\omega)||$, then $H(\omega)$ is the best Chebyshev approximation to $F(\omega)$ from the set of rational functions $\overline{R}_{m,n}$.

Therefore, if we obtain a rational function with the alternation property over a sufficient number of extremal frequencies, we are assured it is optimal. However, this alternation is not necessary; the size of the reference set of the best approximation may be less than $m + n + 2$, and in this case, the best approximation is called *degenerate*. Degeneracy occurs when the leading cosine coefficients of the best approximation, a_m and b_n, are both zero. For more information, see the discussion of the *defect* of the best rational approximation in [**?**] or [**?**].

The progression of the rational Remez algorithm relies on the following key fact. If (i) the set S over which the approximation is performed consists of exactly $n + m + 2$ points and (ii) the best approximation does indeed have $m + n + 2$ extremal frequencies, then the best approximation over S can be found by solving equation (**??**). This is an interpolation problem and its solution is explained below.

The rational Remez algorithm follows the same strategy as the polynomial Remez algorithm:

1. **Initialization:** Select a reference set of $n + m + 2$ points.

2. **Interpolation:** Calculate the best approximation to $F(\omega)$ over this reference set. (Solve the system (**??**)).

3. **Update:** Update the reference set exactly as in the polynomial Remez algorithm. Go back to step 2.

Interpolation Step: Although the system in (**??**) is nonlinear in the coefficients of $H(\omega) = A(\omega)/B(\omega)$, it can be written as a generalized eigenvalue problem [?]: rewrite (**??**) as

$$A(\omega_i) + (-1)^i \delta B(\omega_i) = F(\omega_i) B(\omega_i)$$

where ω_i for $i = 1, ..., m+n+2$ is the current reference set and the unknowns are δ and the coefficients of $A(\omega)$ and $B(\omega)$. $|\delta|$ is called the levelled reference error.

In matrix notation, we have

$$\mathbf{M_1 a} + \delta \mathbf{D_1 M_2 b} = \mathbf{D_2 M_2 b} \tag{3}$$

where

$$\begin{aligned}
\mathbf{a} &= (a_0, ..., a_m)^t &(4)\\
\mathbf{b} &= (1, b_1, ..., b_n)^t &(5)\\
(\mathbf{M_1 a})_i &= A(\omega_i) &(6)\\
(\mathbf{M_2 b})_i &= B(\omega_i) &(7)\\
(\mathbf{D_1})_{i,i} &= (-1)^i &(8)\\
(\mathbf{D_2})_{i,i} &= F(\omega_i) &(9)
\end{aligned}$$

Specifically, $\mathbf{D_1}$ and $\mathbf{D_2}$ are diagonal matrices and

$$\mathbf{M_1} = \begin{bmatrix} 1 & \cos(\omega_1) & \cdots & \cos(m\omega_1) \\ \vdots & & & \vdots \\ 1 & \cos(\omega_L) & \cdots & \cos(m\omega_L) \end{bmatrix}$$

(where $L = m+n+2$) and similarly for $\mathbf{M_2}$. $\mathbf{M_1}$ is a matrix of size $m+n+2$ by $m+1$ and has full rank $m+1$. Therefore, there is a matrix $\mathbf{Q}$ of size $n+1$ by $m+n+2$ with full rank $n+1$ such that $\mathbf{QM_1} = \mathbf{0}$. Applying $\mathbf{Q}$ to (**??**) we eliminate $\mathbf{a}$ and obtain the equation for δ and $\mathbf{b}$

$$\delta \mathbf{Q D_1 M_2 b} = \mathbf{Q D_2 M_2 b}. \tag{10}$$

Once δ and $\mathbf{b}$ are found, $\mathbf{a}$ is found by solving a linear system (see equation **??**). Equation (**??**) is a generalized eigenvalue problem (it is of the form $\mathbf{Ax} = \lambda \mathbf{Bx}$). Notice that the vector $\mathbf{b}$ can be freely scaled without affecting the equality. Therefore, we can scale it so that $b_0 = 1$ assuming

$b_0 \neq 1$. Since there are $n+1$ generalized eigenvalues δ, we must choose an appropriate one. This is straightforward because there will be at most one generalized eigenvalue for which the corresponding denominator $B(\omega)$ is positive over the current reference set [?, ?]. If there is no such value, then the best approximation from $\overline{R}_{m,n}$ to $F(\omega)$ over the reference set is degenerate: it has fewer than $m+n+2$ extremal points. However, even if there is an generalized eigenvalue that gives rise to a denominator positive over the reference set, it may become negative elsewhere on S, the domain of approximation. In either case, the Remez algorithm fails and one must use some corrective measure or an alternative approximation method (see below).

The rational Remez algorithm may fail for two reasons, but in both cases, the failure shows up in the same way: the reference set on some iteration gives rise to no positive denominator. The two reasons the algorithm may fail are:

1. The best approximation from $\overline{R}_{m,n}$ to $F(\omega)$ over S is degenerate. In this case, either the best approximation from $\overline{R}_{m,n}$ over some reference set in the course of the algorithm is degenerate, or the algorithm yields a sequence of approximations that approach degeneracy.

2. Sensitivity to the initial reference set. In this case, the algorithm fails even though the best approximation from $\overline{R}_{m,n}$ is not degenerate.

Unfortunately, it is not possible (to our knowledge) to decide at the time of failure which of these two reasons led to failure. If it is known that the best approximation is degenerate, then the order of the approximation should be reduced.

It is interesting to note that degeneracy of the best approximation over the set S is very rare: for a given function, all intervals on which it has degenerate best approximations form a set of measure zero [?]. For this reason, we assume in this paper that the best approximation is non-degenerate. Near degenerate best approximations are, however, not uncommon. Furthermore, it is the nearly degenerate best approximations that are more computationally difficult to find, for they are sensitive to the initial reference set and unless the usual reference set update procedure is modified, failure of the the rational Remez algorithm for these cases is imminent.

If $E^*_{m,n}$ denotes the Chebyshev error of the best approximation from $\overline{R}_{m,n}$ and if the best approximation from $\overline{R}_{m,n}$ is nearly degenerate, then $E^*_{m-1,n-1}$, the Chebyshev error of the best approximation from $\overline{R}_{m-1,n-1}$, is usually only slightly higher than $E^*_{m,n}$. That is, by reducing the number of poles and zeros both by one, a nearly equivalent approximation can be obtained. Therefore, in the design of optimal recursive filters according to the Chebyshev norm, it is advantageous to reduce the order in this way. For by doing so, the computation required for implementing the filter is reduced while the increase in the Chebyshev error is small. (See example 2 below.) Although a nearly degenerate best approximation may be discarded in preference for a lower order best approximation, the ability to compute nearly degenerate best approximations is nevertheless valuable for the purposes of comparison.

Updating the Reference Set Assuming the algorithm has not failed, the reference set is updated in exactly the same way as in the polynomial Remez algorithm. That is, a new reference set is found such that

1. The current error function on the new reference set alternates sign.
2. The absolute value of the current error function at each point of the new reference set is *at least* as great as the current levelled reference error.
3. The absolute value of the current error function on at least one point of the new reference set must be strictly greater than the current levelled reference error.

In a single point exchange algorithm, one point of the reference set is updated from one iteration to the next. In a multiple point exchange algorithm (which usually converges faster) more than one point is updated, usually the entire reference set. But as long as the three conditions above are satisfied and there is a corresponding positive denominator, the levelled reference error will increase.

Convergence: As in the polynomial Remez algorithm, it can be shown that the levelled reference error $|\delta|$ increases from one iteration to the next as long as the reference set at each iteration gives rise to a positive denominator. Moreover, on each iteration, $|\delta|$ gives a lower bound for the Chebyshev error of the best approximation to $F(\omega)$ over S, for it is itself the Chebyshev error associated with the best approximation over a reference set, a subset of S. That is, $|\delta| < E^*$, where E^* is the Chebyshev error of the best approximation. On each iteration, an upper bound for E^* is given simply by the maximum of the absolute value of the error function, $E(\omega)$. By inspecting at each iteration the upper and lower bound for E^*, it is possible to measure how close the current approximation is to the best approximation. As in the polynomial Remez algorithm, this gives a meaningful stopping criteria.

3 Overcoming Faulty Reference Sets

When no solution to the generalized eigenvalue problem of the interpolation step gives rise to a positive denominator, we suggest perturbing the reference set in a systematic manner.

Suppose that the reference set on some iteration gives rise to a positive denominator (there exists a generalized eigenvector solving **??** that are the coefficients of a positive cosine polynomial). As noted above, it may be the case that the new reference set obtained by updating the current reference set with a multiple (or single) point exchange scheme may fail to give rise to a positive denominator. If this is the case, then the usual rational Remez algorithm that employs the multiple or single point exchange scheme fails.

One way of overcoming this failure is given by the *differential correction* algorithm [**?**, **?**, **?**, **?**]. The differential correction algorithm is a method for calculating best rational Chebyshev approximations by solving a sequence of linear programming problems. It is possible to combine the Remez and differential correction algorithm as is done in [**?**], but because the differential correction algorithm is itself an iterative procedure, we prefer another method for overcoming failure explained as follows.

The single point exchange scheme for updating the reference set is typically carried out by first finding the point, call it ω_{new}, at which $|E(\omega)|$ attains its maximum value and second, by replacing a point in the reference set by ω_{new}. The appropriate point to replace, call it ω_r, is uniquely determined by the conditions listed above for updating the reference set. After updating the reference set by this single point exchange scheme, the reference set is the same as the old one for the exception of one point.

If the reference set obtained by the single point exchange scheme fails to provide a positive denominator, instead of replacing ω_r by ω_{new}, our approach replaces ω_r by $(\omega_r + \omega_{new})/2$. If ω_r and ω_{new} are located on opposite ends of $[0, \pi]$, as occasionally occurs, then $(\omega_r + \omega_{new})/2$ is greater than π and in this case, subtracting π is necessary. If the resulting reference set again fails to provide a positive denominator or if $|E((\omega_r + \omega_{new})/2)| < |\delta|$, then our approach replaces ω_r by $\frac{3}{4}\omega_r + \frac{1}{4}\omega_{new}$. For as long as the new reference fails to provide a positive denominator and an increase in $|\delta|$, the levelled reference error, our approach replaces ω_r by $(1 - \frac{1}{2^k})\omega_r + \frac{1}{2^k}\omega_{new}$. That is, our approach employs successively smaller perturbations to the current reference set.

If no viable reference set is found, then, with respect to the grid density, the new reference point $(1 - \frac{1}{2^k})\omega_r + \frac{1}{2^k}\omega_{new}$ will eventually equal ω_r. In this case, our approach uses another value for ω_{new}. Namely, ω_{new} is taken to be the point at which $|E(\omega)|$ attains its *second* greatest local maximum. With this new value of ω_{new}, our approach carries out the single point exchange again, and subsequently replaces ω_r by $(1 - \frac{1}{2^k})\omega_r + \frac{1}{2^k}\omega_{new}$ for $k = 1, 2, 3, ...$ until a viable reference set is found. Again, if none is found, ω_{new} is taken to be the point at which $|E(\omega)|$ attains its *third* greatest local maximum, and so on.

By testing this sequence of candidate reference set updates, our approach usually finds one that yields a positive denominator and an increase in $|\delta|$. Continuing in this manner usually results in successful convergence to the best approximation. Our observations indicate that when the best approximation is nearly degenerate, then this systematic sequence of perturbations is essential for convergence, for the standard multiple and single point exchange schemes nearly always lead to failure when the best approximation is nearly degenerate.

Sometimes however, before the best approximation is obtained, no perturbation of the reference set by a grid point results in a viable reference set. When this is the case, either the best approximation is actually degenerate, or more likely, more than a perturbation is needed to obtain a reference set from which the Remez algorithm can be made to converge. In our experience, this occurrence can be overcome by moving a reference point from one end of the interval $[0, \pi]$ and inserting it between the two reference points on the other side of the interval.

These observations were collected primarily from experiences with the design of low pass filters, but it is our expectation that the same phenomena are found in general and that the same corrective measures will prove useful. The preceding discussion also assumes that a viable initial reference set has been found. Usually it is not difficult to find an initial reference set giving rise to a positive denominator, although we have not arrived at an entirely robust method for doing so.

4 Constrained Rational Remez Algorithm

In the design of recursive filters, we wish to find a *nonnegative* function approximating the desired magnitude squared frequency response, for we must spectrally factor the approximating function to obtain a stable recursive filter. This constrained approximation is addressed in [**?**] for FIR filter design. Furthermore, the optimality property of the resulting approximations is maintained [**?**]. Here we make appropriate modifications to the rational Remez algorithm.

We impose a constraint on the maximum and minimum values of $H(\omega)$. We call these constraints $u(\omega)$ and $l(\omega)$ for 'upper' and 'lower'. We modify the interpolation step by constructing the rational function interpolating $F(\omega_i) + (-1)^i\delta$, $u(\omega_i)$, or $l(\omega_i)$ at ω_i depending on the error function. The lower constraint is violated at ω_i if

$$F(\omega_i) - |\delta| sgn(F(\omega_i) - H(\omega_i)) < l(\omega) \tag{11}$$

while the upper constraint is violated at ω_i if

$$F(\omega_i) + |\delta| sgn(H(\omega_i) - F(\omega_i)) > u(\omega). \tag{12}$$

The resulting equations are as above, (**??**), but

$$(\mathbf{D_1})_{i,i} = \begin{cases} 0 & \text{if (\textbf{??}) or (\textbf{??}) at } \omega_i \\ (-1)^i & \text{else} \end{cases} \tag{13}$$

$$(\mathbf{D_2})_{i,i} = \begin{cases} u(\omega_i) & \text{if (\textbf{??}) at } \omega_i \\ l(\omega_i) & \text{if (\textbf{??}) at } \omega_i \\ F(\omega_i) & \text{else} \end{cases} \tag{14}$$

As above there is a matrix $\mathbf{Q}$ such that $\mathbf{QM_1} = \mathbf{0}$, and by applying $\mathbf{Q}$ to (**??**) we obtain again a generalized eigenvalue problem. Except for the differences in $\mathbf{D_1}$ and $\mathbf{D_2}$, the interpolation step of the (upper and lower bound) constrained and unconstrained Remez algorithms are the same.

Updating the reference set from one iteration to the next in the constrained Remez algorithm requires some more care than it does in the unconstrained version. For the unconstrained version, it is sufficient to use the value of the error function at its local maxima and minima to choose new reference points. However, for the constrained version, it is necessary to check points of $H(\omega)$ that violate the constraints. While the unconstrained version uses $|H(\omega_i)| - |\delta|$ to select new reference points (this value should be positive), the unconstrained version should use this value at points where the constraint is not violated and one of the values, $l(\omega_i) - H(\omega_i)$ or $H(\omega_i) - u(\omega_i)$, where the upper or lower bound constraint is violated.

In order to obtain non-negative approximations, we simply take $l(\omega)$ to be 0 and we do not use $u(\omega)$. In order to design Elliptic filters, it is necessary to take $u(\omega)$ to be 1.

5 Examples

Example 1 In this example the 13 minimum phase filters with a total number of poles and zeros equal to 12 were designed for an ideal low pass filter with a pass band edge at $1341\pi/2048$ and a stop band edge at $1390\pi/2048$ (so that the band edge is at $2\pi/3$). The total number of grid points used was 2049 for the interval $[0, \pi]$ including the end points and the zero weighted transition band. The pole-zeros plots, the magnitude squared frequency responses and the error functions are shown in figures **??** through **??**.

The Chebyshev error as a function of the number of zeros is listed in table 1 and plotted in figure **??**. As can be seen from this data, the use of two poles significantly reduces the Chebyshev error of the best approximation. It is also interesting to observe that the all pole filter gives a smaller Chebyshev error than does the FIR filter. Of course, for practical purposes of implementation, this does *not* suggest that the all pole filter is preferable to the FIR filter. This is because the phase and the effects of finite precision arithmetic must be taken into account.

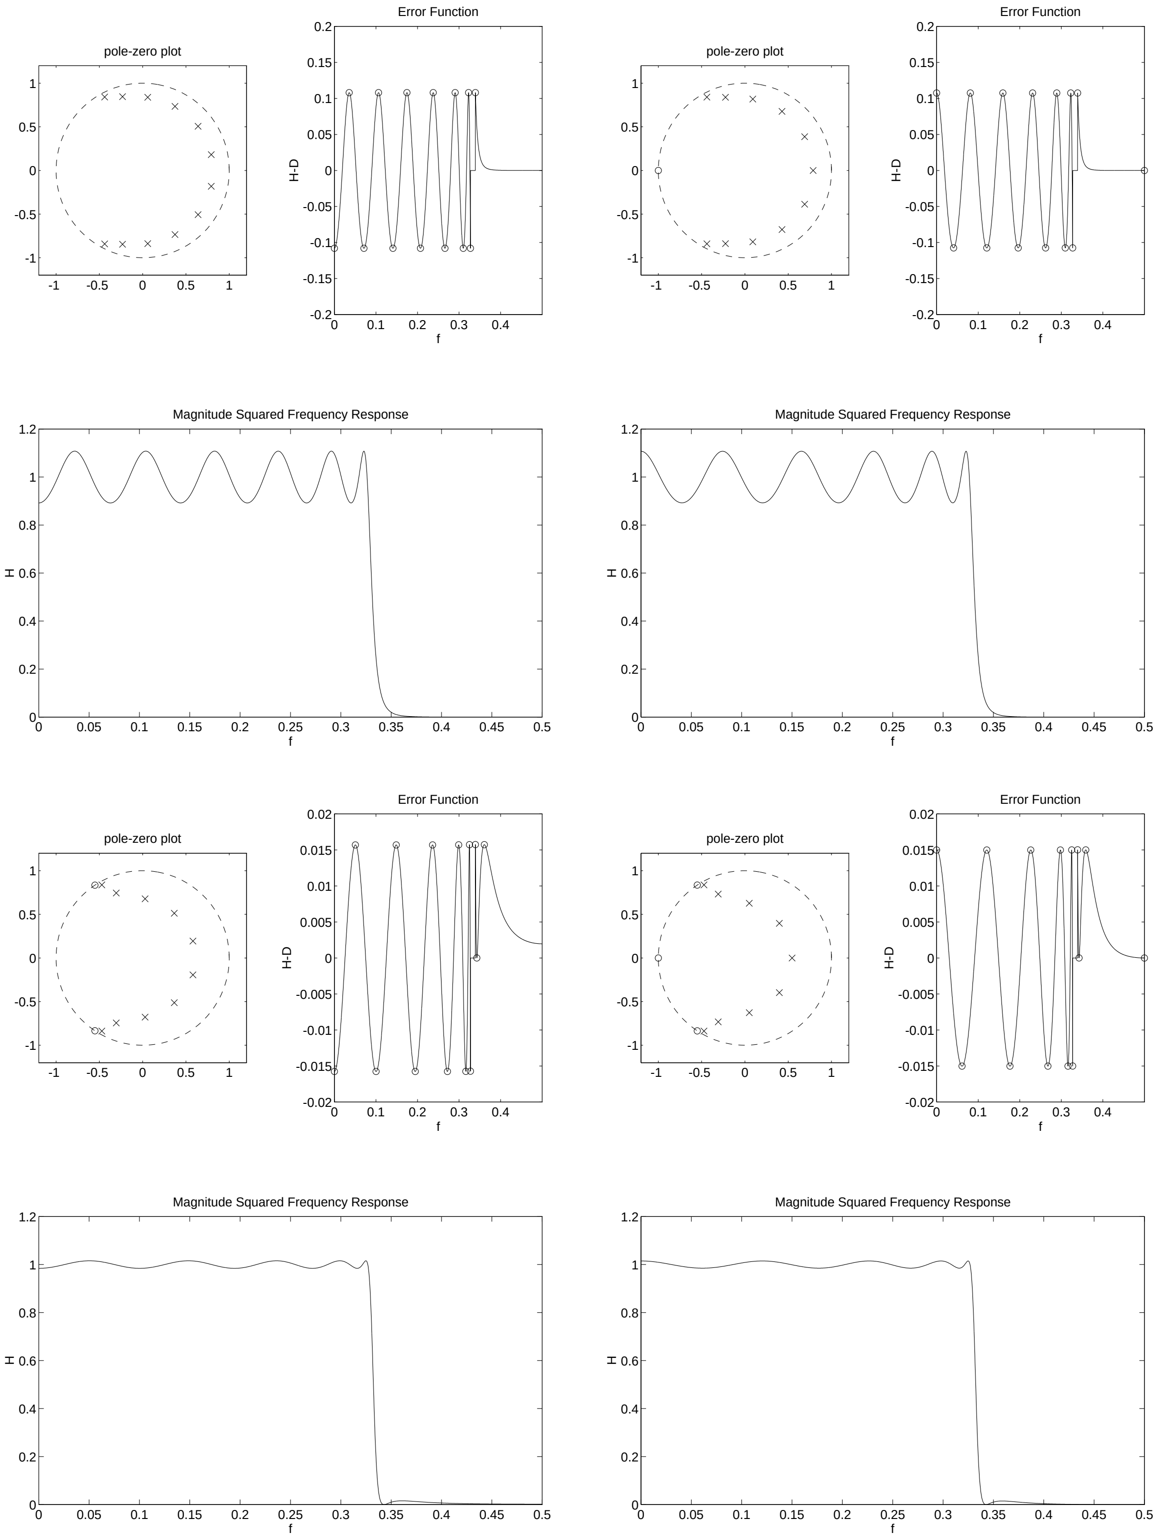

Figure 1: The filters for example 1 having 0,1,2,and 3 zeros.

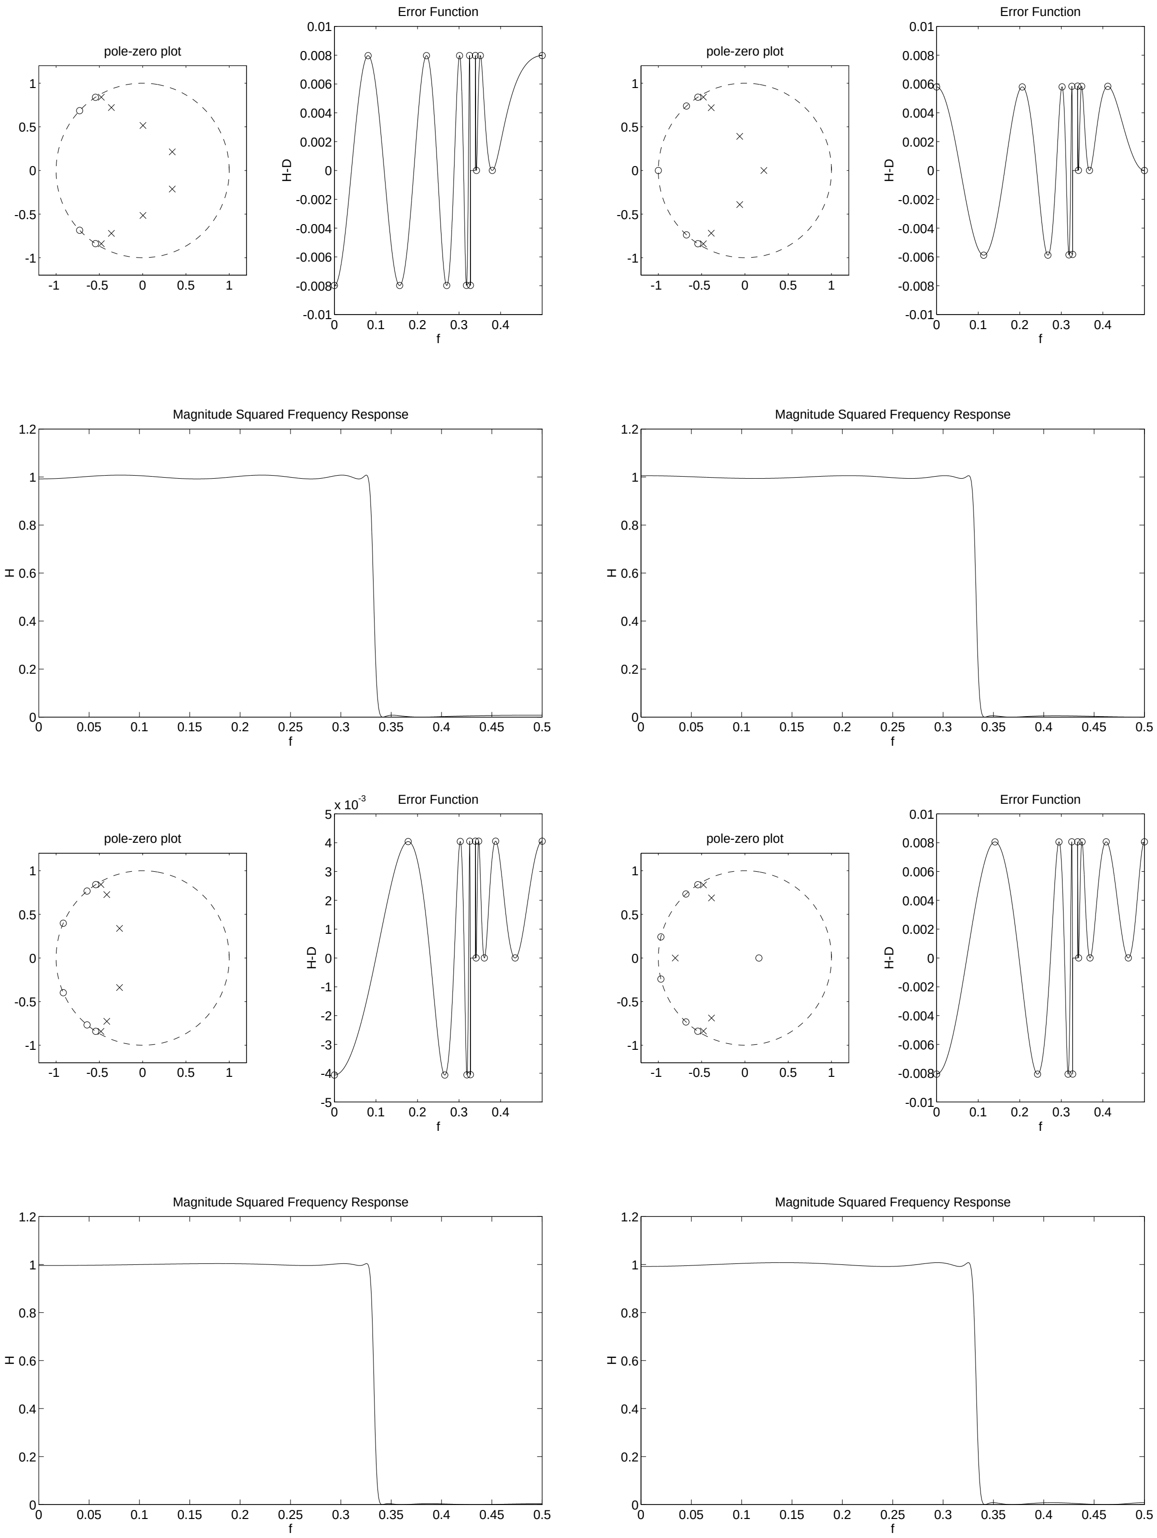

Figure 2: The filters for example 1 having 4,5,6,and 7 zeros.

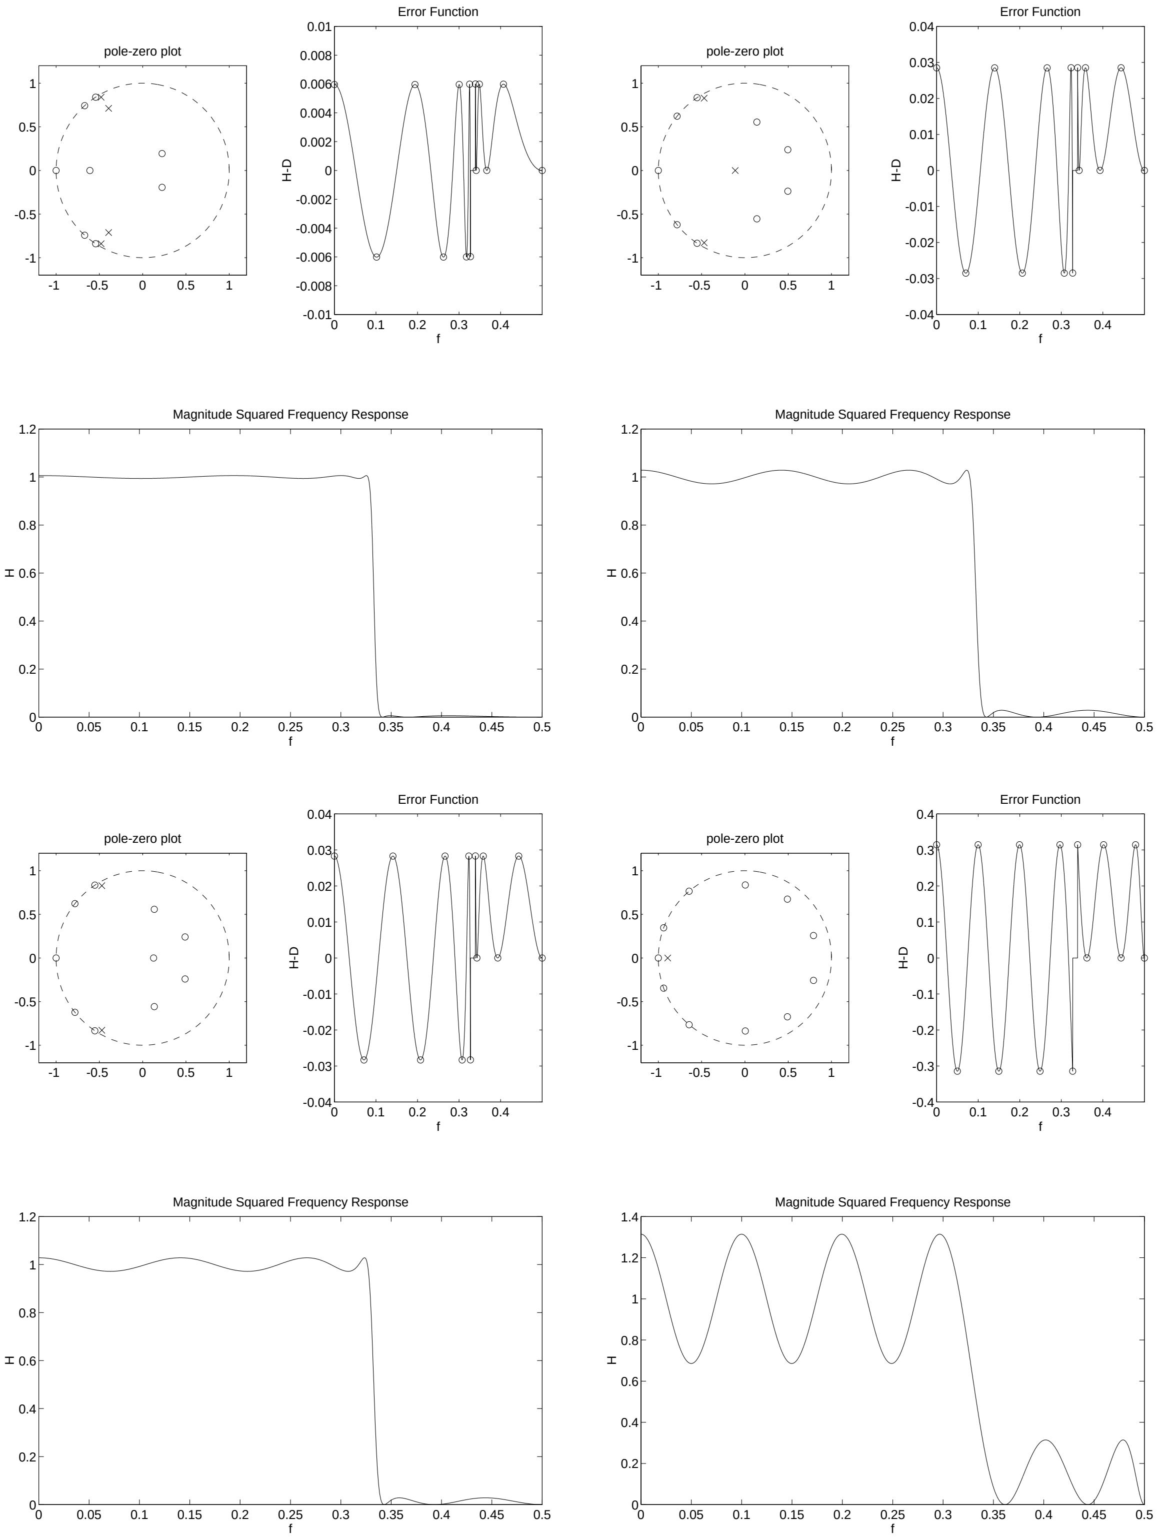

Figure 3: The filters for example 1 having 8,9,10,and 11 zeros.

Table 1: The Chebyshev error of the low pass filters designed for example 1.

number of zeros	$\|\|E\|\|_\infty$
0	0.10796808
1	0.10749376
2	0.01572686
3	0.01500134
4	0.00797507
5	0.00584038
6	0.00405150
7	0.00805965
8	0.00598926
9	0.02850165
10	0.02829758
11	0.31430428
12	0.30123874

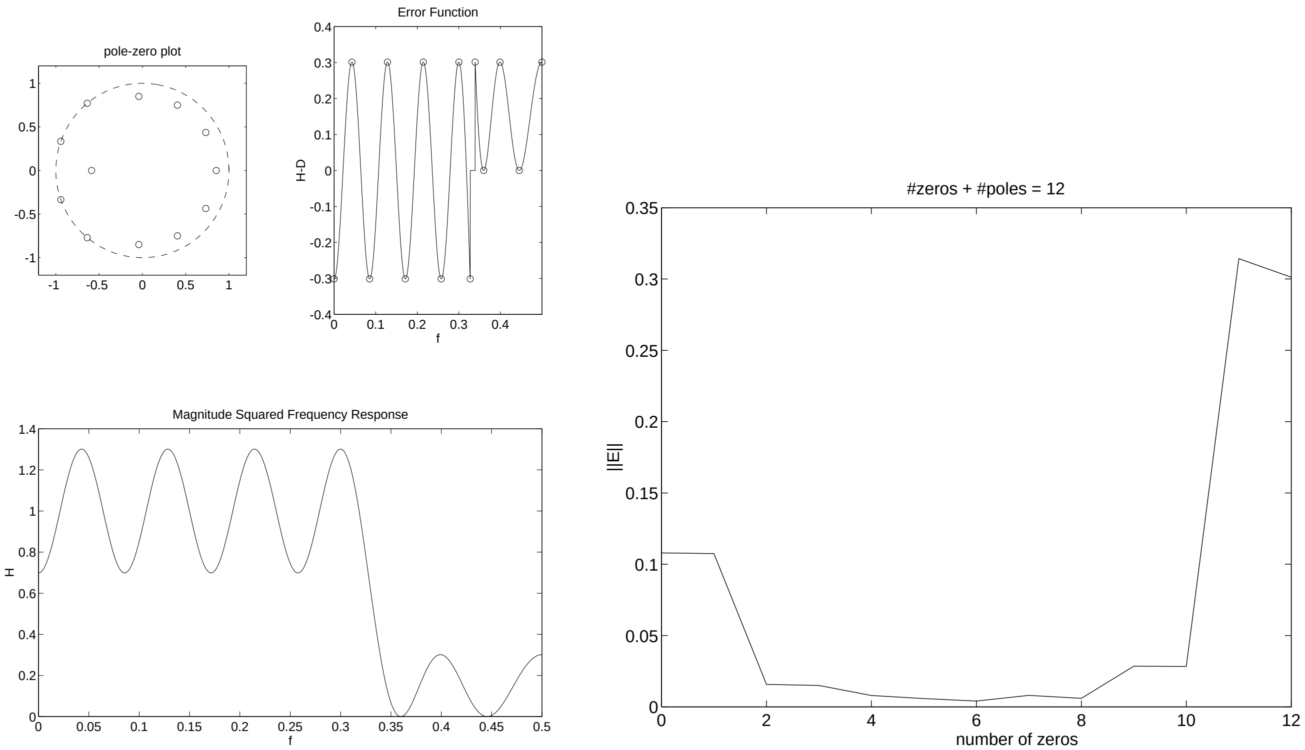

Figure 4: The FIR filter for example 1 and the Chebyshev error as a function of the number of zeros.

Note that for this example, when the number of zeros is greater than 6, the optimal filter possesses zeros lying off the unit circle. For these cases, the optimal filter can not be found with the methods for filter design requiring the zeros to lie on the unit circle.

Example 2 In this example, we design a filter with 8 zeros and 3 poles whose magnitude squared frequency response is nearly degenerate. The ideal frequency response is a low pass filter with a pass band edge at $1426\pi/2048$ and a stop band edge at $1475\pi/2048$. The total number of grid points used was 2049 for the interval $[0, \pi]$ including the end points and the zero weighed transition band. This is an example in which updating the reference set from iteration to iteration requires small perturbations, for the usual exchange methods lead to failure.

The Chebyshev error for the resulting filter was $E^*_{8,3} = 0.040120$. The pole-zero plots, the magnitude squared frequency response and the error function are shown in figure **??**. As can be seen, a pole and a zero almost cancel as is typical for nearly degenerate approximations. Here, the zero is at $z = -1$ and the pole is just inside the unit circle on the real line.

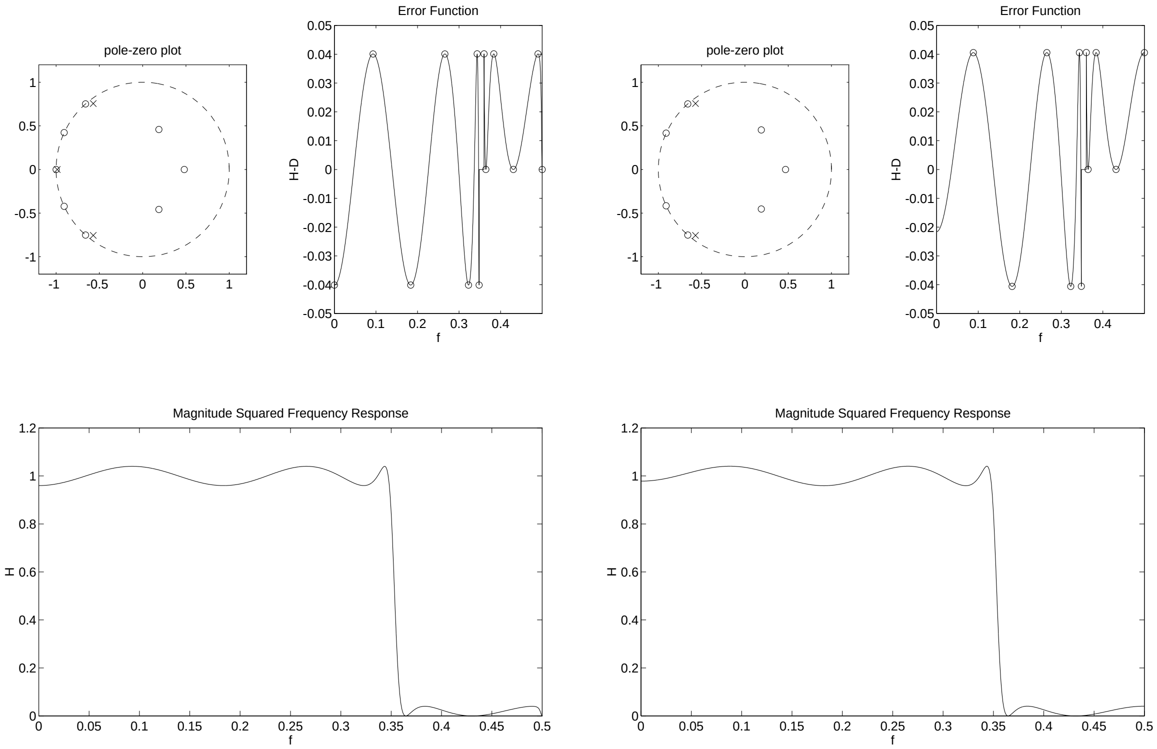

Figure 5: The filters for example 2.

Since a pole and zero almost cancel, it makes sense for practical considerations to decrease the number of poles and zeros by one each. The resulting lower order filter is no longer nearly degenerate and the Chebyshev error is only slightly greater at $E^*_{7,2} = 0.040581$. Notice that the lower order filter has an 'extra' ripple. However, the frequency at which this extra ripple occurs is not an extremal point, for $|E(\omega)|$ does not attain its maximum there.

In general, as the best approximation for a fixed number of poles and zeros becomes more degenerate, the size of the extra ripple in the best approximation of lower order rises to the Chebyshev error. When the best approximation is in fact degenerate, then there is exact pole-zero cancellation and the best approximation of lower order is identical.

If the degree of the approximating function is reduced by reducing only the number of poles by one, then one obtains $E^*_{8,2} = 0.040483$. If the number of zeros is reduced by one, then one gets $E^*_{7,3} = 0.040487$. As expected, $E^*_{8,2}$ and $E^*_{7,3}$ lie between $E^*_{8,3}$ and $E^*_{7,2}$, suggesting that, since $E^*_{8,3} \approx E^*_{7,2}$, the best trade-off between complexity and quality of approximation is given by the filter with 7 zeros and 2 poles. That is, when an approximation is nearly degenerate, reducing both the number of zeros and poles by one generally makes sense.

Example 3 This example illustrates the flexibility of the constrained rational Remez algorithm by designing a filter with a non-conventional frequency response. We use more zeros than poles and a weighting function that is not a constant.

The ideal magnitude frequency response was taken to be a low pass filter with a pass band edge at $2928\pi/2048$ and a stop band edge at $3217\pi/2048$ (so that the band edge is at $.75\pi$). The magnitude frequency response was taken to have a linear increase in the pass band: it was taken to be 0.9 at $\omega = 0$ and 1.1 at the pass band edge, with a linear slope between $\omega = 0$ and the pass band edge. The total number of grid points used was 2049 for the interval $[0, \pi]$ including the end points and the zero weighted transition band. The weighting function $W(\omega)$ was taken to be $1/T(\omega)$ where $T(\omega)$ (the tolerance function) is 1 at $\omega = 0$ and 0.2 at the pass band edge with a linear slope between these two frequencies. $T(\omega)$ was taken to be 0.1 in the stop band.

Figure **??** shows the pole-zero plot, the (weighted) error function, the magnitude squared frequency response, and the magnitude frequency response. Figure **??** shows the un-weighted error function.

Note that the error function shown in figure **??** is as defined above: it is the difference between the realized and desired *squared* magnitudes. For this reason, it is often useful to weight the error more heavily in the stop band, or wherever the ideal response is small, for squaring the magnitude frequency response decreases the error in these regions much more than it does in the pass band. The dashed line in figure **??** is $E^* \cdot T(\omega)$ where E^* is the Chebyshev error of the best approximation.

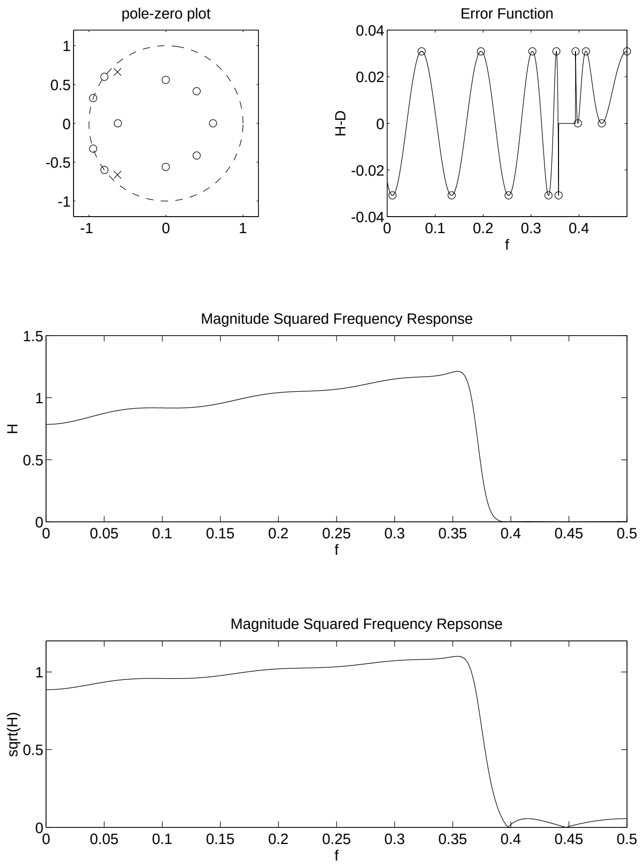

Figure 6: The filter for example 3.

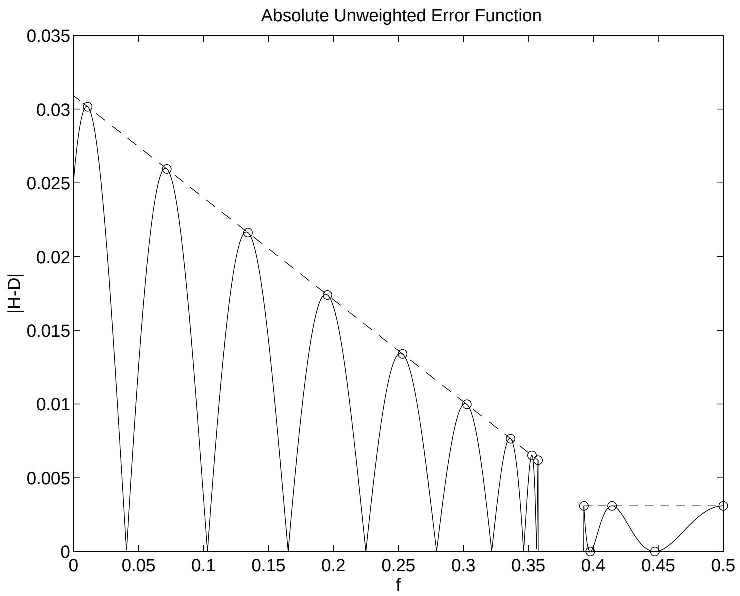

Figure 7: The un-weighted error function for example 3.

6 Summary

We have described a flexible, efficient Remez algorithm for the magnitude squared design of recursive digital filters in the frequency domain. The number of poles and zeros can be chosen arbitrarily and the zeros do not have to lie on the unit circle. This algorithm allows us to design filters with non-conventional frequency responses with arbitrary weighting functions. Moreover, this Remez algorithm can be used to design optimal minimum phase FIR filters and yields Elliptic filters as special cases.

We have addressed three main difficulties in the use of the Remez algorithm for recursive filter design: We use the generalized eigenvalue problem to solve the relevant nonlinear equations of the interpolation step. We impose nonnegativity constraints so that spectral factorization can be employed. Reference set degeneracy is overcome by adjusting the reference set using a sequence of successively smaller perturbations.

Three examples were given illustrating the usefulness of the constrained rational Remez algorithm for the design of recursive filters. The first example illustrated the way in which the Chebyshev error of the optimal filter behaves as a function of the number of zeros when the number of poles and zeros is kept constant. The second example examined a nearly degenerate best approximation and aspects of near degeneracy were discussed. The third example showed the flexibility of the algorithm by using a non-conventional frequency response with a non-constant weighting function.

Acknowledgment

This research has been supported by DARPA under NDSEG fellowship N00014-89J-3204, by Alexander von Humboldt-Stiftung and AFOSR under grant F49620-1-0006, and by BNR.