Design of Recursive Digital Filters in Parallel Form by Linearly Constrained Pole Optimization

Esteban Maestre, Gary P. Scavone, and Julius O. Smith

***Abstract*—We present a technique to iteratively optimize poles of a recursive digital filter in parallel form. Only exposing the poles as the variables to optimize, we employ a linearly constrained gradient descent routine in which the numerical estimation of the error gradient involves first obtaining the zeros by projecting the target response over a basis of responses defined by the pole positions at a given step. Example fits are presented for exponentially decaying white noise and measured violin radiativity filters.**

***Index Terms*—Design of recursive digital filters, optimization, parallel form.**

I. Introduction

RECURSIVE digital filters in parallel form offer a number of advantages. Besides their suitability for parallel implementation, improved numerical precision and parametric control are some of their attractive features. In the context of transfer-function modeling for audio equalization, the design of recursive digital filters over a warped frequency axis has proven to be an effective methodology to reduce the order of a filter while maintaining a desired accuracy over perceptual frequency scales [1]; however, available techniques are often based on minimizing error measures that poorly represent the amplitude compression mechanisms of the human ear [2]. It is difficult to find techniques for automatic design of recursive parallel filters from prescribed responses, and existing workarounds (e.g., [3]–[5]) rely on parallel formulations where parallel sections are constructed from uniform pole distributions or from the poles of direct-form solutions obtained by minimization of magnitude ratio (e.g., linear prediction), l_2-norm (e.g., Steiglitz–McBride iteration), or l_∞-norm [6].

In this letter, we propose an iterative numerical scheme to optimize the design of digital filters in parallel form. At each iteration only pole positions are exposed as the variables to optimize: once they are decided, zeros are constrained to minimize an auxiliary quadratic cost function, resulting in a simple closed-form solution. The positions of the poles are optimized, therefore, iteratively, through gradient descent. The scheme, for which a set of linear constraints is added to ease convergence, offers flexibility in defining the objective error measure, making it possible to minimize a classic l_2-norm or a perceptually motivated *logarithmic spectral distance* (LSD). We study the performance of our technique in modeling randomly generated responses, over different filter orders and two different error measures. Finally, we provide details on how to utilize this technique over a warped frequency axis.

Manuscript received June 9, 2016; revised August 15, 2016; accepted August 23, 2016. Date of publication September 1, 2016; date of current version September 19, 2016. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Balázs Bank.

E. Maestre is with the Computational Acoustic Modeling Lab, McGill University, Montréal, QC H3A 0G4, Canada, and also with the Music Technology Group, Universitat Pompeu Fabra, Barcelona 08002, Spain (e-mail: esteban@ccrma.stanford.edu).

G. P. Scavone is with the Computational Acoustic Modeling Lab, McGill University, Montréal, QC H3A 0G4, Canada (e-mail: gary.scavone@mcgill.ca).

J. O. Smith is with the Center for Computer Research in Music and Acoustics, Stanford University, Stanford, CA 94305 USA (e-mail: jos@ccrma.stanford.edu).

Digital Object Identifier 10.1109/LSP.2016.2605626

II. Design of Parallel Filters From a Fixed Pole Set

The z-domain general expression of the parallel filter model under consideration is

$$\hat{H}(z) = g + \sum_{n=1}^{N} R_n(z) + \sum_{m=1}^{M} P_m(z) \tag{1}$$

where the scalar term g is real and each of the parallel terms

$$R_n(z) = \frac{c_n}{1 - r_n z^{-1}}$$

$$P_m(z) = \frac{b_{m,0} + b_{m,1} z^{-1}}{(1 - p_m z^{-1})(1 - p_m^* z^{-1})} \tag{2}$$

respectively corresponds to a real pole section and a complex-conjugate pole pair section, both with real numerator coefficients. The order of the filter is $L = N + 2M$. We can decompose each term $P_m(z)$ in (1) as

$$P_m(z) = b_{m,0}\, Q_m(z) + b_{m,1}\, z^{-1} Q_m(z) \tag{3}$$

where

$$Q_m(z) = \frac{1}{(1 - p_m z^{-1})(1 - p_m^* z^{-1})} \tag{4}$$

corresponds to a two-pole real resonator. Given N real poles $r_1 \cdots r_n \cdots r_N$ and M complex poles $p_1 \cdots p_m \cdots p_M$, it is straightforward to formulate a problem that is linear in the numerator coefficients [4]. Let vector $\mathbf{h} = [h_1 \cdots h_k \cdots h_K]^T$ contain K samples of the target frequency response $H(e^{j\omega})$ at frequencies $0 \leq \omega_k < \pi$, i.e., $h_k = H(e^{j\omega_k})$. Likewise, let vector $\mathbf{r}_n = [r_{n,1} \cdots r_{n,k} \cdots r_{n,K}]^T$ contain the normalized sampled frequency-response $r_{n,k} = R_n(e^{j\omega_k})$ with $c_n = 1$, and similarly let $\mathbf{q}_m^0 = [q_{m,1}^0 \cdots q_{m,k}^0 \cdots q_{m,K}^0]^T$ with $q_{m,k}^0 = Q_m(e^{j\omega_k})$, and $\mathbf{q}_m^1 = [q_{m,1}^1 \cdots q_{m,k}^1 \cdots q_{m,K}^1]^T$ with $q_{m,k}^1 = e^{-j\omega_k} Q_m(e^{j\omega_k})$. Next, let $\mathbf{A}$ be the $K \times N + 2M + 1$ matrix of basis vectors constructed as

$$\mathbf{A} = [\mathbf{1}, \mathbf{r}_1 \cdots \mathbf{r}_N, \mathbf{q}_1^0 \cdots \mathbf{q}_M^0, \mathbf{q}_1^1 \cdots \mathbf{q}_M^1] \tag{5}$$

where $\mathbf{1} = [1 \cdots 1]^T$ denotes the length K vector of all ones. Let vector $\mathbf{d}$ contain the projection coefficients arranged as

$$\mathbf{d} = [g, c_1 \cdots c_N, b_{1,0} \cdots b_{M,0}, b_{1,1} \cdots b_{M,1}]^T. \tag{6}$$

Given the pole positions r_n and p_m, we form the corresponding matrix $\mathbf{A}$, and the numerator coefficients $\mathbf{d}$ can be found by solving the least-squares projection problem

$$\underset{\mathbf{d}}{\text{minimize}} \ \|\mathbf{A}\mathbf{d} - \mathbf{h}\|^2. \tag{7}$$

A related unconstrained algorithm was proposed for direct-form recursive filters by Burrus and Parks [7], [8], and particularly related is the variation of Prony's method which computes the numerator coefficients of the transfer function by minimizing the "solution error" using fixed denominator coefficients designed previously by minimizing "equation error" [8, pp. 226–228].

III. Optimization of Pole Positions

We first parameterize an initial set of N real stable pole positions and a set of M complex-conjugate stable pole pair positions as follows. We group the initial N real poles into a set of U positive poles and a set of V negative poles, and represent each uth and vth pole by its respective radius parameter $s_u = -\log(1 - |r_u|)$ and $s_v = -\log(1 - |r_v|)$. Then, we represent each mth complex pole pair in terms of its angle parameter $w_m = |\angle p_m|$ and its radius parameter $s_m = -\log(1 - |p_m|)$. This leads to four parameter sets: a set $\mathbf{u}$ of radius parameter values corresponding to the U positive real poles, a set $\mathbf{v}$ of radius parameter values corresponding to the V real negative poles, and sets $\mathbf{w}$ and $\mathbf{s}$ of angle and radius parameter values, respectively, corresponding to the M complex-conjugate pole pairs. Finally, a key step is to sort the pole parameter sets so that linear constraints can be defined in a straightforward manner to ensure that the arrangement of poles in the unit circle is preserved during optimization. Besides reducing the search space, this constraint avoids potential singularities at pole frequency crossings. Elements in sets $\mathbf{u}$ and $\mathbf{v}$ are respectively sorted by ascending radius parameters s_u and s_v, while elements in sets $\mathbf{w}$ and $\mathbf{s}$ are jointly sorted as pairs (each pair corresponding to a complex-conjugate pole) by ascending angle parameter w_m.

We pose our pole optimization problem as

$$\begin{aligned} \underset{\mathbf{u},\mathbf{v},\mathbf{w},\mathbf{s}}{\text{minimize}} \quad & \varepsilon(H, \hat{H}) \\ \text{subject to} \quad & \mathbf{C} \end{aligned} \tag{8}$$

where $\varepsilon(H, \hat{H})$ is an error measure between the target frequency response and model frequency response, and $\mathbf{C}$ is a set of linear constraints employed to ensure feasibility and to aid convergence. To solve this problem, we can resort to numeric optimization via a gradient descent routine that makes use of local quadratic approximations of $\varepsilon(H, \hat{H})$ by successive numerical evaluations [9]. Constraints $\mathbf{C}$ are defined as follows. First, feasibility is ensured by

$$\begin{aligned} &0 \le s_u && \forall u \in [1, U] \\ &0 \le s_v && \forall v \in [1, V] \\ &0 \le s_m && \forall m \in [1, M] \\ &0 \le w_m < \pi && \forall m \in [1, M]. \end{aligned} \tag{9}$$

Second, to aid convergence, we constrain the pole sequence order in sets $\mathbf{u}$, $\mathbf{v}$, and $\mathbf{w}$ to be respected, so that all real poles appear in ascending radius parameter and complex poles appear in ascending angle parameter. This is expressed by

$$\begin{aligned} s_{u-1} &< s_u < s_{u+1} && \forall u \in [2, U-1] \\ s_{v-1} &< s_v < s_{v+1} && \forall v \in [2, V-1] \\ w_{m-1} &< w_m < w_{m+1} && \forall m \in [2, M-1]. \end{aligned} \tag{10}$$

Moreover, assuming that initialization provides an already trusted first solution, we can bound the search to a region around the initial pole positions. This can be expressed via the additional inequalities

$$\begin{aligned} s_u^- &< s_u < s_u^+ && \forall u \in [1, U] \\ s_v^- &< s_v < s_v^+ && \forall v \in [1, V] \\ w_m^- &< w_m < w_m^+ && \forall m \in [1, M] \\ s_m^- &< s_m < s_m^+ && \forall m \in [1, M] \end{aligned} \tag{11}$$

where '$-$' and '$+$' superscripts are used to respectively indicate static lower and upper bounds defined during initialization.

For the optimization routine to successfully approximate the error gradient, we must supply a procedure to evaluate the error function $\varepsilon(H, \hat{H}|_j)$ at step j as a function of the model parameters $\mathbf{u}|_j$, $\mathbf{v}|_j$, $\mathbf{w}|_j$, and $\mathbf{s}|_j$ at step j. This is carried out in two steps. First, from the pole positions at iteration j, numerator coefficients are obtained by solving the linear problem as detailed in Section II. Then, the resulting model frequency response $\hat{H}(e^{j\omega})|_j$ is evaluated at K frequencies $0 \le \omega_k < \pi$, with samples arranged in a vector $\hat{\mathbf{h}}|_j = [\hat{h}_1|_j \cdots \hat{h}_k|_j \cdots \hat{h}_K|_j]^T$ that is used to compute an arbitrarily defined error measure with respect to the target response vector $\mathbf{h}$. For example, the error measure can be defined to be the l_2-norm

$$\varepsilon_2(H, \hat{H}|_j) = \left\| \hat{\mathbf{h}}|_j - \mathbf{h} \right\|^2 \tag{12}$$

or a perceptually motivated [2], [10] LSD

$$\varepsilon_{\text{LSD}}(H, \hat{H}|_j) = \sum_{k=1}^{K} \left| \log_{10} \left| \hat{h}_k|_j \right| - \log_{10} \left| h_k \right| \right|. \tag{13}$$

.

In Fig. 1, we compare optimized designs for both error measures. We display the magnitude absolute error $\epsilon_2(\omega_k) = \left| |\hat{h}(\omega_k)| - |h(\omega_k)| \right|$ and the log-magnitude absolute error $\epsilon_{\text{LSD}}(\omega_k) = \left| \log_{10} |\hat{h}(\omega_k)| - \log_{10} |h(\omega_k)| \right|$ of each optimized design, with initial pole positions defined by the roots of the denominator of the proper transfer function as obtained via Prony's method. It is clear how the ε_2-optimized design does a better job at matching the peaks of the magnitude response near $\pi/2$ and $3\pi/4$, at the expense of a higher logarithmic error around the dips near $3\pi/4$ and π.

IV. Performance Analysis

We studied the proposed technique by designing parallel filters from randomly generated target responses, using three different pole initialization techniques and filter orders $L = \{20,\ 30, \ldots, 100\}$. (i) *Uniform* ("Unif"): we uniformly distribute $(L-2)/2$ complex-conjugate pole pairs over the unit disc and place two poles (one positive and one negative) on the real axis, all with equal radius. (ii) *Prony's method* ("Prony"): we initialize L poles by computing the roots of the L-order denominator polynomial of a proper transfer function whose coefficients are obtained by fitting the target impulse response

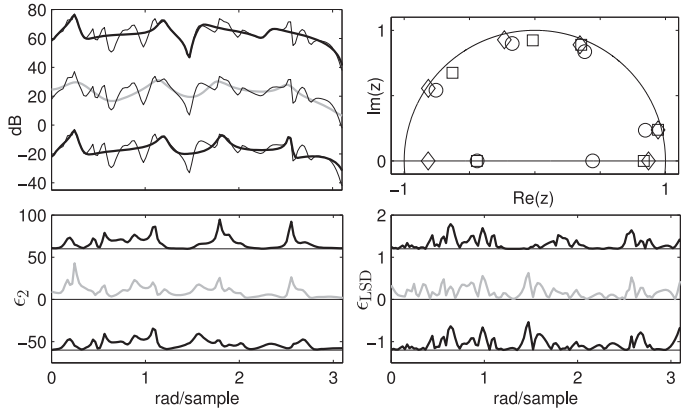

Fig. 1. Example design of order $L = 10$, comparing the ε_{LSD}-*optimized* (top curves) and the ε_2-*optimized* (bottom curves) solutions for a randomly generated response of length $K = 2^8$, and initial pole positions provided by Prony's method (middle curves). Pole positions (positive imaginary axis only) are depicted in the top-right plot: '$\circ$', '$\square$', and '$\diamond$' symbols respectively used for initial, ε_{LSD}-*optimized*, and ε_2-*optimized* solutions.

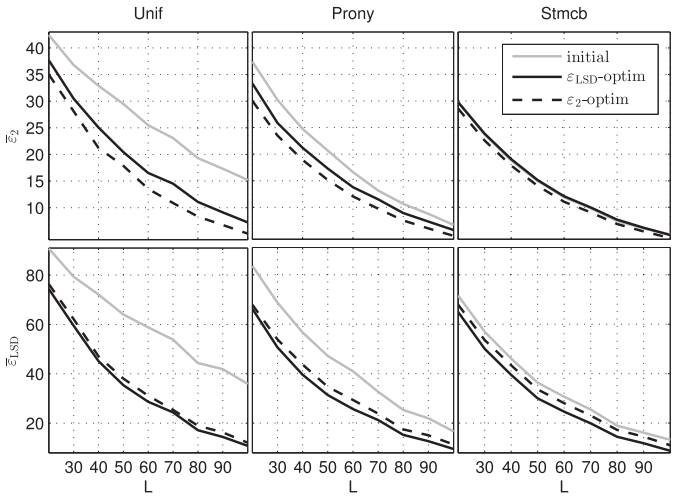

Fig. 2. Comparative analysis results for different values of L, showing the error measures averaged over 50 trials on randomly generated target responses.

via Matlab's `prony` routine. (iii) *Steiglitz–McBride* ("Stmcb"): analogously, we initialize L poles by fitting the target frequency response via Matlab's `stmcb` routine for a maximum of 300 iterations. For each filter order L, 50 target responses were constructed from exponentially decaying white noise sequences of length $K = 2^{10}$. For each target response and initialization technique, we designed three filters: a *direct* design from initial pole positions (no optimization, see Section II), a ε_2-*optimized* design (12), and a ε_{LSD}-*optimized* design (13), and computed both error measures for each. For optimization, we used Matlab's `fmincon` interior-point gradient descent routine as the basis for the implementation of our method, with a maximum of 300 iterations. For all optimization trials we employed feasibility constraints (9) and pole order constraints (10). Fig. 2 shows the values of ε_2 and ε_{LSD} of the direct and optimized designs, averaged over 50 trials for each order L. In all cases, the optimization is able to reduce both error measures, except for the ε_{LSD}-*optimized* solution from initial poles provided by "Stmcb" initialization, which could not reduce the ε_2 error.

V. Warped Frequency Design

To solve (7) on a warped frequency axis, one could use nonuniform frequency sampling. However, in a gradient descent context, it may be desirable to also remap pole positions into a warped unit circle so that their distribution in the (warped) search space matches a warped frequency sampling density used for error estimation. This matching can be achieved by using the bilinear conformal map [1] to reformulate (7) as follows. Mapping of $R_n(z)$ in (2) to the warped ζ-plane yields

$$R_n(\zeta) = c_n X_n(\zeta) \tag{14}$$

with

$$X_n(\zeta) = \chi_n \frac{1 + \lambda\zeta^{-1}}{1 - \alpha_n\zeta^{-1}} \tag{15}$$

where $0 < \lambda < 1$ is the warping parameter, the pole α_n is related to its linear-frequency counterpart a_n by

$$a_n = \frac{\alpha_n + \lambda}{1 + \lambda\alpha_n} \tag{16}$$

and $\chi_n = (1 - \lambda a_n)^{-1}$. Likewise, mapping of $P_m(z)$ in (3) to the ζ-plane yields

$$P_m(\zeta) = b_{m,0} P_m^0(\zeta) + b_{m,1} P_m^1(\zeta) \tag{17}$$

with

$$\begin{aligned} P_m^0(\zeta) &= \gamma_m (1 + \lambda\zeta^{-1})^2 Y_m(\zeta) \\ P_m^1(\zeta) &= \gamma_m \lambda (1 + \lambda\zeta^{-1})(1 + \lambda^{-1}\zeta^{-1}) Y_m(\zeta) \end{aligned} \tag{18}$$

where

$$Y_m(\zeta) = \frac{1}{(1 - \rho_m\zeta^{-1})(1 - \rho_m^*\zeta^{-1})} \tag{19}$$

is a two-pole resonator defined by a pair of complex-conjugate poles ρ_m and ρ_m^* that are related to their linear-frequency counterparts p_m and p_m^* by

$$p_m = \frac{\rho_m + \lambda}{1 + \lambda\rho_m} \tag{20}$$

and $\gamma_m = (1 - 2\lambda|p_m|\cos\angle p_m + \lambda^2|p_m|^2)^{-1}$. With this formulation, it is possible to solve for linear-frequency numerator coefficients in the warped frequency axis as follows. We first construct a target warped-frequency response vector $\boldsymbol{\eta}$ by complex resampling of $\mathbf{h}$ at the corresponding frequencies in the ζ-plane. Then, we evaluate the frequency responses of $X_n(\zeta)$, $P_m^0(\zeta)$, and $P_m^1(\zeta)$ and arrange samples into vectors $\mathbf{x}_n$, $\mathbf{p}_m^0$, and $\mathbf{p}_m^1$ respectively so that a matrix $\boldsymbol{\Lambda}$ can be constructed in an analogous way to matrix $\mathbf{A}$ in (5). Finally, solving the least-squares projection problem

$$\underset{\mathbf{d}}{\text{minimize}} \; \|\boldsymbol{\Lambda}\mathbf{d} - \boldsymbol{\eta}\|^2 \tag{21}$$

yields linear-frequency numerator coefficients in $\mathbf{d}$ as they appear in (6). Successively solving (21) allows preforming the complete optimization process on a warped frequency scale as follows: (i) warp target response, initial poles, and constraints; (ii) optimize warped poles; (iii) dewarp optimized poles and use last solution to (21) as numerator coefficients. In Fig. 3, we provide an illustrative example. We randomly generated a target response of length $K = 2^9$, and initialized poles by uniformly

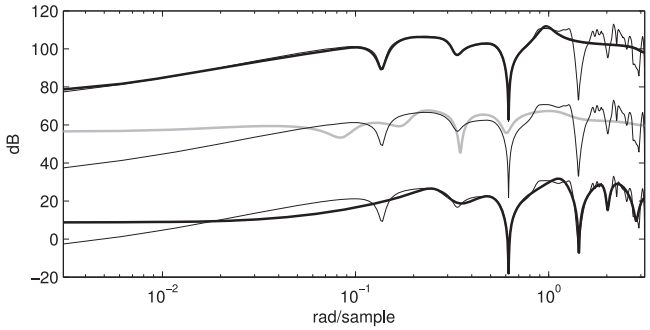

Fig. 3. Example $\varepsilon_{\mathrm{LSD}}$-*optimized* solutions ($L = 12$) for a warped frequency design ($\lambda = 0.75$, top) and a linear frequency design ($\lambda = 0$, bottom), compared to the initial solution (middle).

distributing 12 poles on the ζ-plane ($\lambda = 0.75$), all with equal radius. We then optimized ($\varepsilon_{\mathrm{LSD}}$ criterion) for the cases $\lambda = 0.75$ and $\lambda = 0$. It can be observed that linear-frequency optimization caused initial poles to move into the high frequencies, while warped-frequency optimization ensured higher resolution in the low end.

Fig. 4 depicts two pairs of warped-frequency design examples obtained from fitting a violin radiation efficiency measurement, displayed in $250 < f < 9600$ Hz. Each pair consists of an initial design and its corresponding ε_2-*optimized* design. The two initializations are obtained similarly as detailed in Section IV. (i) *Warped Prony's method* ("wProny"): warp target response; obtain L poles via Matlab's `prony` routine; dewarp poles. (ii) *Warped Steiglitz–McBride* ("wStmcb"): warp target response; obtain L poles via Matlab's `stmcb` routine for 300 iterations; dewarp poles. Numerator coefficients of initial and optimized designs are found by solving (21). For "wProny", pole optimization reduced ε_2 from 1.472 to 1.223; for "wStmcb", ε_2 decreased from 0.923 to 0.872.

VI. Conclusion

In this letter, we have presented a numerical scheme to optimize the poles of recursive digital filters in parallel form. The scheme relies on a series of iterations exposing only the positions of the poles as the variables to optimize, while zeros are not free variables: once the pole locations are decided, the numerator coefficients are calculated to minimize an auxiliary quadratic cost function, resulting in a simple closed-form solution. The positions of the poles are optimized iteratively through gradient descent. Feasibility constraints are added to ease convergence. We have shown how the technique is able to consistently improve the modeling accuracy for different pole initialization methods, filter orders, and two error measures: a classic l_2-norm and a LSD. We have provided details on how to adapt the technique to work over a warped frequency axis via bilinear transform.

Besides imposing a formulation that guarantees stability and provides flexibility for defining the objective error measure, this technique allows optimizing the coefficients of a filter directly in parallel form, providing the improved numerical properties and dynamic-range benefits of parallel form for filters that are most naturally implemented in parallel, such as modal resonator banks. Beyond providing a comparison between this technique

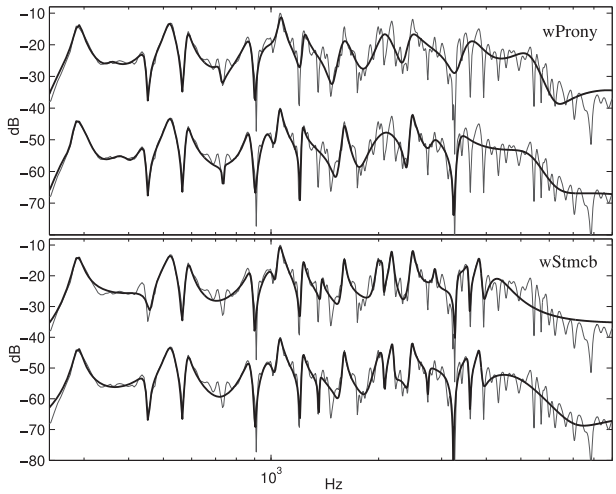

Fig. 4. Warped designs ($L = 40$, $\lambda = 0.8$) from a violin radiativity profile. From top: wProny, ε_2-*optimized* wProny, wStmcb, ε_2-*optimized* wStmcb.

and trusted previous methods, the presented performance analysis was carried out to test the reliability of the method. An attractive potential of this technique lies in its flexibility when initializing or constraining the desired frequency distribution of poles.

Convergence to a satisfactory solution is greatly influenced by initialization and constraints, and we have observed occasional spurious divergences that can be tempered via tightening the constraints. In that regard, an appropriate upper bound for radius parameters may help in cases where excessively high pole selectivity induces worse numerical evaluations of the error function. Future studies of different solvers and other initialization techniques (e.g., spectral peak picking) may help to give a better understanding of convergence properties and to further exploit the potential applications of the method.

References

[1] J. O. Smith and J. S. Abel, "Bark and ERB bilinear transforms," *IEEE Trans. Speech Audio Process.*, vol. 7, no. 6, pp. 697–708, Nov. 1999.
[2] M. Karjalainen, "Sound quality measurements of audio systems based on models of auditory perception," in *Proc. IEEE Conf. Acoust., Speech, Signal Process.*, 1984, pp. 132–135.
[3] M. Karjalainen and T. Paatero, "High-resolution parametric modeling of string instrument sounds," in *Proc. IEEE Eur. Signal Process. Conf.*, 2005, pp. 1–4.
[4] B. Bank and G. Ramos, "Improved pole positioning for parallel filters based on spectral smoothing and multiband warping," *IEEE Signal Process. Lett.*, vol. 18, no. 5, pp. 299–302, May 2013.
[5] J. Rämö, V. Välimäki, and B. Bank, "High-precision parallel graphic equalizer," *IEEE Trans. Audio, Speech, Language Process.*, vol. 22, no. 12, pp. 1894–1903, Dec. 2014.
[6] A. Deczky, "Synthesis of recursive digital filters using the minimum p-error criterion," *IEEE Trans. Audio Electroacoust.*, vol. 20, no. 4, pp. 257–263, Oct. 1972.
[7] C. S. Burrus and T. W. Parks, "Time domain design of recursive digital filters," *IEEE Trans. Audio Electroacoust.*, vol. 18, no. 2, pp. 137–141, Jun. 1970.
[8] T. W. Parks and C. S. Burrus, *Digital Filter Design*. New York, NY, USA: Wiley, Jun. 1987, contains FORTRAN software listings.
[9] J. Nocedal and S. J. Wright, *Numerical Optimization*. New York, NY, USA: Springer, 2006.
[10] S. Wu and W. Putnam, "Minimum perceptual spectral distance FIR filter design," in *Proc. IEEE Conf. Acoust., Speech, Signal Process.*, 1997, pp. 447–450.