forum acusticum 2026
8-12 September · Graz, Austria

DISCRETE-TIME IIR FILTER DESIGN FOR RADIAL FILTERS BY NUMERICALLY OPTIMISED POLE/ZERO PLACEMENT

Frank Schultz[1], Nara Hahn[2], Sascha Spors[1]

[1] *Institute of Communications Engineering, University of Rostock, Germany*
[2] *Institute of Sound and Vibration Research, University of Southampton, UK*

This paper is a revised version of the authors' submitted manuscript that was peer-reviewed by at least two anonymous reviewers.

Abstract

For a certain mode in spherical wave field expansions, a radial filter characterises the radially dependent diffraction phenomenon due to a spherical obstacle. Recently, the band-limited impulse invariance method (BLIIM) was proposed. It allows a more precise design of discrete-time filters for improved numerical simulation of such diffraction phenomena. However, for frequencies close to the Nyquist frequency and/or for high modal orders, BLIIM is numerically not straightforward or not feasible. Then, one possible workaround is an optimisation-based infinite impulse response filter design with the continuous-time frequency response as target function, and solving for the coefficients of the z-domain transfer function.

An optimisation problem is exemplarily discussed for the radial filter occurring in the velocity-to-pressure spherical exterior expansion. This high-pass-shaped radial filter exhibits a steep slope and a rippled pass band for a high modal order. The numerical experiments show that precise magnitude and phase responses of such radial filters can be achieved, requiring only some informed heuristics for the filter design parameters. The approach can be used where BLIIM is not feasible.

Keywords: *radial filter, frequency sampling method, band-limited impulse invariance method, numerical filter design*

1. Introduction

The recently introduced band-limited impulse invariance method (BLIIM) [1] allows for improved design of discrete-time impulse responses from continuous-time recursive models. In [2, 3] this was utilised for radial filters that act in spherical wave field expansions.

The propagation problem from [3] aims on modelling spatial aliasing, for which a proper spherical expansion requires high modal orders for large temporal frequencies. The involved high-pass filters exhibit a steep slope and a rippled pass band, and can be analytically given as Laplace transfer function and subsequently as partial fraction expansion, which requires numerical computation of the poles and the residues. This is numerically straightforward for low modal order n. In the experiments of [3], finding the poles with numpy's `roots()` is numerically incorrect for modal orders $n > 25$. Consequently, the designed least squares (LS) BLIIM filters considerably degrade in quality. Hence, for the wave propagation simulation from a dodecahedron source [3, Fig. 4], only the modes $n \leq 24$ were considered to properly model the spatial aliasing in the 4 kHz third-octave band using 32 kHz sampling frequency.

The present paper considers a recursive filter design by numerical optimisation for higher order modes, where LS-BLIIM is not available due to current numerical issues. Section 2 presents the chosen method, followed by Section 3 discussing the performance of filter designs. Section 4 gives concluding remarks.

2. Method

Outgoing wave propagation from a pulsating sphere with radius R and a given radial velocity profile is considered. The spherical wave field extrapolation to the sound pressure along a sphere with radius r in free-field requires a high-pass-like radial filter for each mode $n = 0 \dots N$ up to an appropriate expansion order N. The frequency response of such a radial filter for the n-th mode is given as [3–5]

$$H_n(\omega, r, R) = -\mathrm{j}\frac{h_n^{(2)}(\frac{\omega}{c}r)}{h_n^{'(2)}(\frac{\omega}{c}R)}, \quad r > R \tag{1}$$

using $\mathrm{e}^{+\mathrm{j}\omega t}$ time convention, n-th order spherical Hankel function of second kind $h_n^{(2)}(\cdot)$, its derivative $h_n^{'(2)}(\cdot)$ [6, Ch. 10.47, 10.51], speed of sound c, angular frequency ω, time t and imaginary unit j. According to

Corresponding author: *frank.schultz@uni-rostock.de.*

paper [3], the Eq. (1) can also be given as

$$H_n(s,r,R) = \frac{R}{r} \mathrm{e}^{-s\cdot(\frac{r-R}{c})} \left(1 - \hat{H}_n^{\mathrm{LP}}(s,r,R)\right) \quad (2)$$

in the Laplace domain ($\mathrm{j}\omega \to s$). In [3], $\hat{H}_n^{\mathrm{LP}}(s,r,R)$ is explicitly stated as a partial fraction expansion representation, which is not needed here, but from which it was deduced that the transfer function $\hat{H}_n^{\mathrm{LP}}(s,r,R)$ is a minimum-phase low-pass filter with n zeros, $n+1$ poles and 1st order integrator characteristics for $\omega \to \infty$, cf. [3, Fig. 1 (right)]. The high-pass filter

$$\hat{H}_n^{\mathrm{HP}}(s,r,R) = \left(1 - \hat{H}_n^{\mathrm{LP}}(s,r,R)\right) \quad (3)$$
$$= H_n(s,r,R)\,\frac{r}{R}\,\mathrm{e}^{+s\cdot(\frac{r-R}{c})} \quad (4)$$

of our interest here is consequently minimum-phase with $n+1$ zeros and $n+1$ poles (the same as of the low-pass) and 1st order differentiator characteristics for $\omega \to 0$ rad/s. The $6(n+1)$ dB/octave high-pass slope determines the n-dependent phase excess in the transition band, cf. [3, Fig. 1 (left)].

Aiming at a discrete-time, recursive filter with infinite impulse response (IIR), that shall closely match the continuous-time impulse/frequency response, the above mentioned LS-BLIIM filter design was utilised in [3] for this dedicated radial filter. In the present paper, we aim for an optimisation-based filter design, such that Eq. (4) closely matches a discrete-time high-pass filter encoded in the z-domain transfer function

$$\hat{H}_n^{\mathrm{HP}}(z) = \Big(\sum_{\mu=0}^{N_z} b_\mu z^{-\mu}\Big) \Big/ \Big(\sum_{\nu=0}^{N_p} a_\nu z^{-\nu}\Big), \quad (5)$$

with chosen fixed coefficient $a_0 = 1$. For the sake of brevity, not all of the dependencies $\hat{H}_n^{\mathrm{HP}}(z,r,R)$, $b_\mu(r,R,n)$ and $a_\nu(r,R,n)$ are labelled, but it is worth noting them. Depending on the number of poles N_p and zeros N_z and their respective location, the digital domain transfer function $\hat{H}_n^{\mathrm{HP}}(z)$ represents either a recursive or a non-recursive system.

A vast number of optimisation approaches for the design of IIR filters exists, cf. e.g. [7–11]. Instead of comparing them, the paper addresses specific design details for the given radial filter problem. For that, we opted for the filter design method [12] with detailed derivations given in [13, Ch. 5]. This is a Gauss-Newton based LS solver to minimise the magnitude of the complex-valued error frequency response. The method designs pure non-recursive as well as recursive filters, and for the latter controls the filter's stability by constraining the maximum allowed pole radius. The algorithm from [12] inspired recent approaches, e.g. [14] essentially realised a minimax problem with this pole constraint. Also, [12] competes with other recent approaches, e.g. cf. the experiments and comparisons in [15, Sec. IV] and the comments on the approach in [16, p.1726]. Hence, [12] is here considered as a powerful non-commercial up-to-date solver.

The optimisation problem with the sampling frequency f_s in Hz is as follows: The error frequency response, weighted with $W \in \mathbb{R} > 0$ (user defined),

$$E(\omega_i) = \sqrt{W(\omega_i)} \cdot \left(\hat{H}_n^{\mathrm{HP}}(s=\mathrm{j}\omega_i)\cdot \mathrm{e}^{-\mathrm{j}\frac{\omega_i}{f_s}\cdot m} - \hat{H}_n^{\mathrm{HP}}(z=\mathrm{e}^{\mathrm{j}\frac{\omega_i}{f_s}})\right) \quad (6)$$

is defined at discrete angular frequencies ω_i by appropriate frequency axis sampling (user defined). The linear phase shift $\mathrm{e}^{-\mathrm{j}\frac{\omega_i}{f_s}\cdot m}$ adds a constant group delay of m/f_s seconds (m, f_s user defined). The calculation of the high-pass target frequency response $\hat{H}_n^{\mathrm{HP}}(s=\mathrm{j}\omega_i)$ with Eq. (4) at specified angular frequencies ω_i relies on the numerically valid computation of the spherical Hankel functions that appear in Eq. (1).

Stacking complex-valued error values, as well as the filter coefficients to be optimised, into columns

$$\boldsymbol{e} = [E(\omega_1),\, E(\omega_2),\, E(\omega_3), \ldots]^{\mathrm{T}} \quad (7)$$
$$\boldsymbol{b} = [b_0,\, b_1, \ldots,\, b_{N_z}]^{\mathrm{T}}, \quad \boldsymbol{a} = [a_1, \ldots,\, a_{N_p}]^{\mathrm{T}}$$

and denoting the ν-th pole of the searched discrete-time high-pass system $\hat{H}_n^{\mathrm{HP}}(z)$ as $z_{\infty,\nu}$, the filter design presented in [12], [13, Ch. 5] asks for

$$\min_{\boldsymbol{b},\boldsymbol{a}} \boldsymbol{e}^{\mathrm{H}}\boldsymbol{e} \quad \text{s.t.} \quad |z_{\infty,\nu\in\{1\ldots N_p\}}| \leq \rho_{\max}, \quad (8)$$

with stable, causal filters requiring $0 \leq \rho_{\max} < 1$ and the maximum allowed pole radius $\rho_{\max}$ being a user-defined constraint. Note that the case $N_z < N_p$ is not considered in the present paper. Further note that the filter design algorithm effectively handles $N_z > N_p$ as

$$\hat{H}_n^{\mathrm{HP}}(z) = \Big(\sum_{\mu=0}^{N_z} b_\mu z^{-\mu}\Big) \Big/ \Big(\sum_{\nu=0}^{N_z} a_\nu z^{-\nu}\Big) \quad (9)$$

with $a_\nu = 0$ for $N_p + 1 \leq \nu \leq N_z$. This yields the intended N_p optimisable poles and additionally $N_z - N_p$ non-optimisable poles being fixed in the z-plane's origin. By doing this, the choice $N_p = 0$ yields a pure non-recursive system, i.e. a finite impulse response (FIR) filter.

The sample delay amount $m > 0$ introduced in Eq. (6) allows for quasi-non-causal components in the discrete-time IR. Analogous to Eq. (3), the z-transforms of the discrete-time systems and their IRs over discrete-time index $k \in \mathbb{Z}$ can be straightforwardly linked

$$\hat{H}_n^{\mathrm{HP}}(z) = z^{-m} - \hat{H}_n^{\mathrm{LP}}(z) \;\bullet\!\!-\!\!\circ \quad (10)$$
$$\hat{h}_n^{\mathrm{HP}}(k) = \frac{\sin\big(\pi(k-m)\big)}{\pi(k-m)} - \hat{h}_n^{\mathrm{LP}}(k) \quad (11)$$

when $m \in \mathbb{N}_0$. For $m \in \mathbb{R}$ the Eq. (10) does not hold. The fractional part of the acoustic delay in Eq. (2)

$$\tau = \frac{r-R}{c} \to m = \big(\tau f_s - \lfloor \tau f_s \rfloor\big) + \text{pre-delay} \quad (12)$$

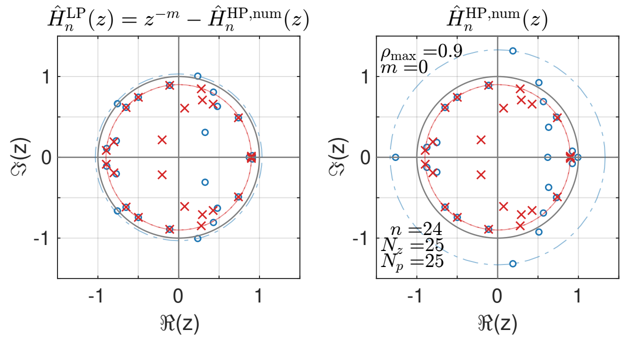

(a) Pole-zero plot of numerically designed high-pass (right) and the corresponding low-pass (left) via Eq. (10). The red dotted circle indicates the user defined maximum pole radius, the red circle the designed maximum pole radius, and the blue dashed circle the maximum zero radius, respectively.

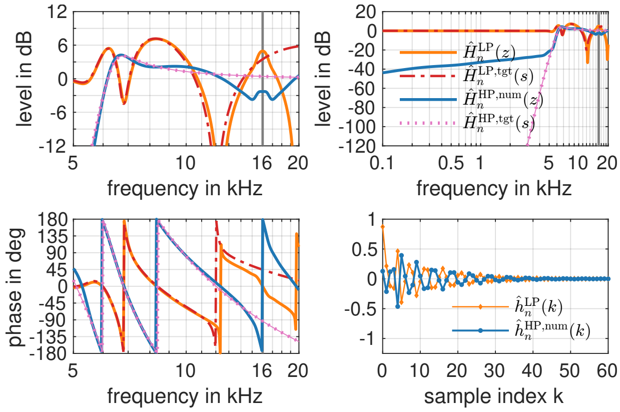

(b) Magnitude / phase / impulse responses for target high-pass ('tgt') vs. numerically designed high-pass ('num'). The corresponding low-passes via Eq. (3), Eq. (10) are shown as well. All phase frequency responses are compensated for the pre-delay m, hence the more pleasing minimum-phase is depicted.

Figure 1: Example I: modal order $n = 24$, $f_s = 32\,\mathrm{kHz}$, $c = 343\,\mathrm{m/s}$, $R = 0.2\,\mathrm{m}$, $r = 150 \cdot \frac{c}{f_s} + R = 1.807\,812\,5\,\mathrm{m}$, $\rho_{\max} = 0.9$, $N_p = 25$, $N_z = 25$, unit error weighting $W = 1$, pre-delay $m = 0$ samples.

can be conveniently added to the filter design of $\hat{H}_n^{\mathrm{HP}}(z) \bullet\!\!-\!\!\circ\, \hat{h}_n^{\mathrm{HP}}(k)$. For the experiments below, the radii r and R are chosen such that $(\tau f_s - \lfloor \tau f_s \rfloor) = 0$. The fractional sample delay case is thus not addressed here in favour of exclusively discussing the minimum-phase part of the filter Eq. (2).

3. Numerical Experiments

The Matlab reference code given in [13] needs the obsolete quadratic program solver `qp()` (Matlab 5) replaced with `quadprog()` (we used Matlab 2023b and 2025b) and few obvious adaptions for the inequality constraints' Lagrange multipliers and for the exit flag. Our Python port of this Matlab reference code is available in the repository[1] dedicated to this paper. Unless otherwise stated, all default parameters of the reference code are retained for the experiments below.

We discuss the performance of the filter design

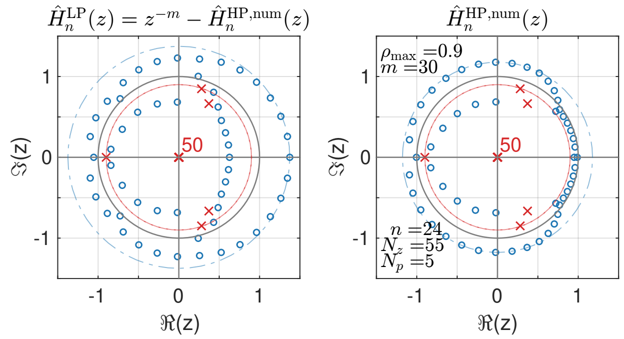

(a) Cf. Fig. 1a. Right: many zeros in the right z-plane close at the unit circle are responsible for the steep transition band. The few poles support the smooth pass band with few ripples. Note that the high-pass (right) and the complementary low-pass (left) have the same poles; this holds for all examples.

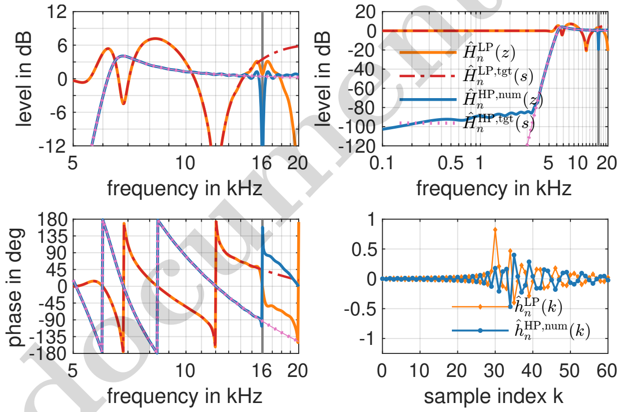

(b) Cf. Fig. 1b.

Figure 2: Example II: modal order $n = 24$, $f_s = 32\,\mathrm{kHz}$, $c = 343\,\mathrm{m/s}$, $R = 0.2\,\mathrm{m}$, $r = 150 \cdot \frac{c}{f_s} + R = 1.807\,812\,5\,\mathrm{m}$, $\rho_{\max} = 0.9$, and parameters different to example I: $N_p = 5$, $N_z = 55$, pre-delay $m = 30$ samples, error weighting $W(f < 6\,\mathrm{kHz}) = 1000$, $W(f > 6\,\mathrm{kHz}) = 1$.

method with five exemplarily chosen examples denoted with Roman numerals I to V, which are also included in the repository. First, poor designs due to inappropriately chosen user parameters (I and III, respectively) are contrasted with useful designs due to reasonable parameters (II and IV, respectively). The filter in example I is somewhat over-complex, whereas the one in example III can be considered under-complex. Second, example V compares the proposed filter design with two other design methods.

The solver is always provided with a log-spaced frequency resolution $[\omega_{i=1} = 2\pi, \ldots, \omega_{i=512} = 2\pi \frac{f_s}{2}]$ rad/s. A sampling frequency $f_s = 32$ kHz is chosen.

It is worth recalling that the continuous-time high-pass $\hat{H}_n^{\mathrm{HP}}(s)$ exhibits $n+1$ poles, $n+1$ zeros and a high-pass slope of $6(n+1)$ dB/octave merging to a differentiating 6dB/octave slope for $\omega \to 0$ rad/s. The latter is not observable in the result graphics due to chosen plot limits. For higher order modes this differentiating slope occurs with already high level attenuation and is likely practically irrelevant anyhow.

[1] `www.github.com/spatialaudio/fa2026_radial_iir_ls`

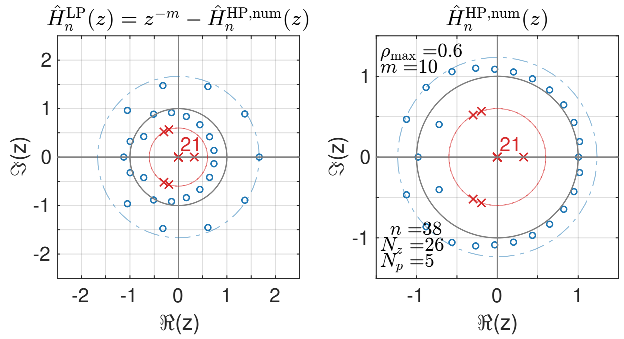

(a) Cf. Fig. 1a.

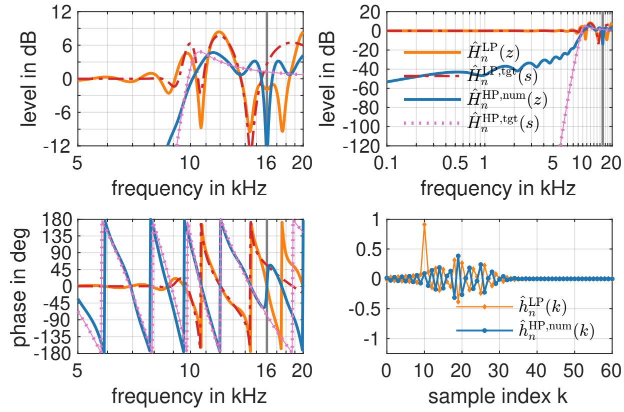

(b) Cf. Fig. 1b.

Figure 3: Example III: modal order $n = 38$, $f_s = 32\,\mathrm{kHz}$, $c = 343\,\mathrm{m/s}$, $R = 0.2\,\mathrm{m}$, $r = 100 \cdot \frac{c}{f_s} + R = 1.271\,875\,\mathrm{m}$, $\rho_{\max} = 0.6$, $N_p = 5$, $N_z = 26$, unit error weighting $W = 1$, pre-delay $m = 10$ samples.

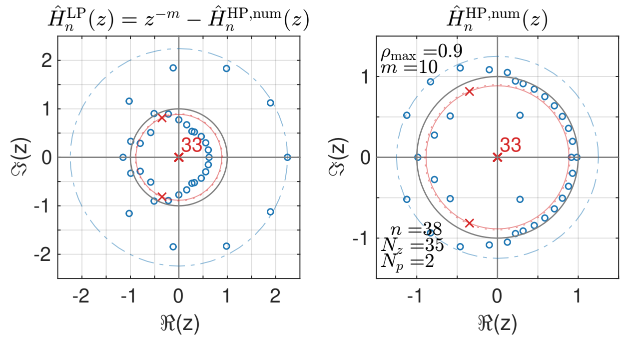

(a) Cf. Fig. 1a.

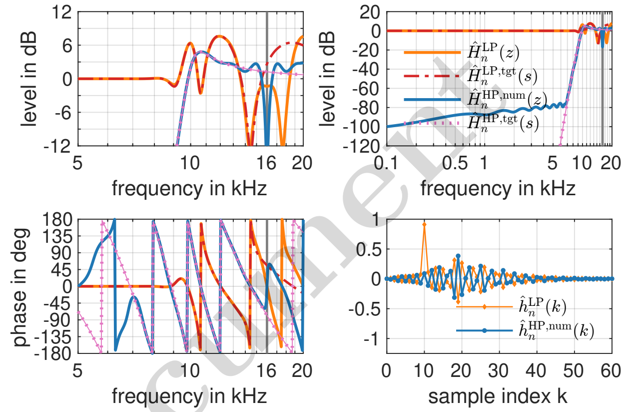

(b) Cf. Fig. 1b.

Figure 4: Example IV: modal order $n = 38$, $f_s = 32\,\mathrm{kHz}$, $c = 343\,\mathrm{m/s}$, $R = 0.2\,\mathrm{m}$, $r = 100 \cdot \frac{c}{f_s} + R = 1.271\,875\,\mathrm{m}$, pre-delay $m = 30$ samples, and parameters different to example III: $\rho_{\max} = 0.9$, $N_p = 2$, $N_z = 35$, error weighting $W(f < 8\,\mathrm{kHz}) = 5000$, $W(f > 8\,\mathrm{kHz}) = 1$.

Example I (cf. Fig. 1) aims at a discrete-time filter $\hat{H}_n^{\mathrm{HP}}(z)$ with $N_p = n+1$ and $N_z = n+1$ for an exemplarily chosen modal order $n = 24$. With a minimum-phase discrete-time filter in mind, the pre-delay is set to $m = 0$ samples. The frequency dependent weight is set to unity for all frequencies and the maximum pole radius is set to 0.9. The pole-zero plot for the designed $\hat{H}_n^{\mathrm{HP}}(z)$ in Fig. 1a (right) reveals that the filter is not minimum-phase, as no optimisation constraint ensures this. However, the phase plot for $\angle \hat{H}_n^{\mathrm{HP}}(z)$ in Fig. 1b (bottom/left, thick/blue) indicates that the filter design fits well to the minimum-phase characteristics of the continuous-time filter in the pass band (dotted/purple), considerably deviating only for $f < 5\,\mathrm{kHz}$ and for $f > 15\,\mathrm{kHz}$. The Bode diagram shows that the target slope of the high-pass filter (dotted/purple) is poorly followed by the designed filter (thick/blue), cf. Fig. 1b (top/right). The slopes deviate considerably for frequencies $f < 5.5\,\mathrm{kHz}$; the designed filter achieves poor DC attenuation. Further, the pass band of the designed high-pass filter does not match well with that of the target, cf. Fig. 1b (top/left). Fig. 1a (right) reveals that some poles in the left z-plane are almost cancelled by zeros and that many other poles are optimised close to $z = 1$, unintentionally heavily supporting instead of attenuating the DC. The effort in using this high amount of poles is unprofitable. These deductions match well with the statements [13, p.175] for denominator degrees above eight to ten. The radial filters here need only few poles to create few ripples/resonances for a comparably small pass band close to the Nyquist frequency.

Example II (cf. Fig. 2) aims for an improved filter. The maximum pole radius is kept at 0.9. A (rather exaggerated) pre-delay of $m = 30$ samples is allowed, setting the number of zeros to $N_z = n+1+m = 55$. The number of poles is considerably decreased to $N_p = 5$. A frequency dependent weighting is introduced: The error in the frequency band $f < 6\,\mathrm{kHz}$ (stop and transition band) is penalised 30 dB more than the frequency band $f > 6\,\mathrm{kHz}$ (pass band). These choices yield a designed high-pass filter with improved performance. The pre-delay and the increased amount of zeros allows the solver to place many zeros close to the unit circle for the stop band region (< 4 kHz), cf. Fig. 2a (right). The two complex-conjugate pole pairs in the right z-plane support the pass band ripples. One pole pair exhibits the maximum allowed pole radius, the other is slightly less resonant. The fifth pole must be real and interacts with a zero almost exactly at $z = -1$, both supporting a resonance at $f_s/2$. This yields the notch at $f_s/2$ in the frequency response of $\hat{H}_n^{\mathrm{HP}}(z)$, but helps the pass band matching the target up to highest frequencies, cf. Fig. 2b (top/left). The steep slope of the transition band (150 dB/octave) is matching well between 4 and 6 kHz, cf. Fig. 2b (top/right). The target and designed slopes deviate for $f < 4\,\mathrm{kHz}$, but an attenuation of > 80 dB for $f < 2.5\,\mathrm{kHz}$ is achieved. This is likely a reasonable performance for practical applications. The designed phase matches the target phase almost exactly for the important frequency band $f > 5\,\mathrm{kHz}$, cf. Fig. 2b (bottom/left). It only considerably deviates close to $f_s/2$ due to the phase excess of the designed notch filter. The IR in Fig. 2b (bottom/right) indicates post-ringing that is similar to example I (due to same maximum pole radius $\rho_{\max} = 0.9$), but also the here allowed pre-ringing due to the pre-delay of $m = 30$ samples is observable.

Example III considers the high modal order $n = 38$ (at which BLIIM is not feasible). The target high-pass exhibits a very steep slope of 234 dB/octave. The cut-off frequency is at about 10 kHz for the chosen wave numbers. A comparably conservative maximum pole radius 0.6 is chosen for demonstration purpose. The pre-delay is set to $m = 10$ samples. The comparably low pole order $N_p = 5$ from example II is retained. The number of zeros is lowered to $N_z = 26$. Unit error weighting is applied. The resulting performance of this filter design is depicted in Fig. 3. Two complex-conjugate pole pairs in the z-plane (here all optimised to $|z_{\infty,\nu}| = \rho_{\max} = 0.6$) support the pass band's ripple characteristics, cf. Fig. 3a (right). Because these four poles are not resonant enough, the pass band matches poorly, Fig. 3b (top/left). For the fifth (real-valued) pole the solver opted for a weak DC-supporting pole at $z \approx +0.326$, hence not helping for the pass band. Another critical aspect is the comparably small amount of zeros for the stop and transition band. In addition to the unit error weighting this leads to a poor match of the slope both in magnitude Fig. 3b (top/right) and in phase Fig. 3b (bottom/left). The IR exhibits a comparably fast decay in the post-ringing part due to the not-so-resonant poles, cf. Fig. 3b (bottom/right). In summary, this filter is neither useful for its transition band nor for its pass band characteristics.

Example IV aims to fix these deficiencies by trading off the accuracy of the pass band ripples vs. that of the transition slope. Based on the pole placement from example III, it is assumed that just one, but then more resonant, complex-conjugate pole pair yields an improved filter design. Thus, $N_p = 2$ and $\rho_{\max} = 0.9$ is chosen. Additionally, as this helped in example II, the error weighting is chosen such that the stop and transition band is penalised 37 dB more than the pass band. The number of zeros is increased to $N_z = 35$, potentially allowing more zeros for the transition band. Note that the chosen $N_z = 35$ is slightly less than the modal order. It was observed that a choice far below the theoretical amount of $N_z = n+1$ generally results in poor designs, and for our experiments a reasonably robust rule of thumb is $N_z = n + 1 + m$ for integer pre-delay m. The performance of this filter design is shown in Fig. 4. The pole-zero plot Fig. 4a (right) shows a complex-conjugate pole pair at the maximum allowed pole radius, which supports the pass band reasonably well up to 13 kHz, cf. Fig. 4b (top/left). The slope is matching well up to about 80 dB attenuation, cf. Fig. 4b (top/right), confirmed by the additional zeros close to the unit circle in the right z-plane in Fig. 4a (right). The phase matches the target well between 7 and 15 kHz, Fig. 4b (bottom/left). Compared to example III, the IR of example IV has a longer post-ringing part, due to the chosen more resonant single pole pair, cf. Fig. 4b (bottom/right). Dedicated for $n = 38$, $R = 0.2\,\mathrm{m}$ and $c = 343\,\mathrm{m/s}$, these filter design parameters could be used for approximately all $r > 2R$. Then, the stop band attenuation is larger than 70 dB and the pass band level error is less than about 1.6 dB. For $R < r < 2R$ dedicated design parameters are needed, due the filter characteristics of the very close near field.

Example V compares the proposed numerical filter design method ('num') with the windowed 'IDFT' approach and with the 'LS-BLIIM'. The results are shown in Fig. 5. Again the modal order $n = 24$ is chosen, but with another radius r and thus another minimum-phase and another pass band than in the examples I and II. For a fair comparison, all filter designs use a $m = 10$ sample pre-delay. The LS-BLIIM filter precisely corresponds to example [3, Fig. 2]. The IDFT-based filter uses a 1024 DFT length from which the first 128 samples represent the FIR. The pre-/post-ringing parts of this FIR are windowed by asymmetrical Kaiser-Bessel fade-in/out windows. The IDFT approach cannot be improved by order of magnitude, hence always achieving the poorest performance for the stop band attenuation in this comparison. The LS-BLIIM performs with a very high attenuation in the stop band. Generally, the performance of the numerical approach can be straightforwardly tuned for poorer slope vs. better pass band, and vice versa. In the shown example we opted for a stop band performance between the IDFT and the LS-BLIIM. This comes with a comparably large pass band ripple, but in return delivers the largest pass band close to $f_s/2$, whereas LS-BLIIM exhibits a considerable low-pass characteristics inherent to its design algorithm.

4. Conclusion

The paper discussed a numerical design for high-pass shaped radial filters. For high modal order these high-pass filters exhibit very steep slopes close to the Nyquist frequency and thus leave only a small, rippled pass band. Thus it was expected and confirmed in this paper, that discrete-time filters modelling this phenomenon, require only few poles (to support the small bandwidth of the pass band and its few resonances) but many zeros (to support the steep slope of the transition band). While the continuous-time filters are minimum-phase, the discrete-time filters are not necessarily, but only the phase response of the

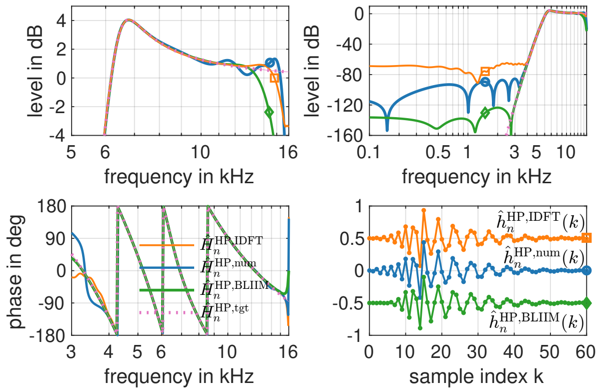

Figure 5: Example V: Comparison of the here discussed filter design (blue, i.e. the middle curve in right plots, 'num') with a windowed IDFT (orange, i.e. the top curve in right plots) and with the LS-BLIIM from [3] (green, i.e. the bottom curve in right plots). General parameters: modal order $n = 24$, $f_s = 32\,\mathrm{kHz}$, $c = 343\,\mathrm{m/s}$, $R = 0.2\,\mathrm{m}$, $r = 200 \cdot \frac{c}{f_s} + R = 2.343\,75\,\mathrm{m}$, pre-delay $m = 10$ samples. IDFT parameters: $N_{\mathrm{fft}} = 1024$, windowed to a 128 samples FIR, 11 fade-in samples, 92 fade-out samples, Kaiser-Bessel $\beta = 5$. Num parameters: $N_p = 4$, $N_z = 35$, $\rho_{\max} = 0.9$, $W(f < 7\,\mathrm{kHz}) = 10^5$, $W(f > 7\,\mathrm{kHz}) = 1$, update step size $\alpha = 0.25$ (not the default value 0.5). BLIIM parameters: 21 FIR coefficients, 27 log-spaced control frequencies [62 . . . 16k] Hz. All phase frequency responses are compensated for the pre-delay m, hence the more pleasing minimum-phase is depicted.

transition and pass band must actually match. The chosen least squares solver with maximum pole radius constraint proved to be a robust choice for the mentioned demands. Some informed heuristics is required for optimum number of poles and zeros and other user definable parameters, but these can be found individually for a specific modal order and the radiating sphere's radius, and then work well for a comparably large wave number range. The proposed numerical design is particularly meant for high modal orders, where the band-limited impulse invariance design is not (yet) available due to numerical limitations in finding the required denominator roots of the Laplace transfer function. A more precise root solver might help and is left for future work. The here proposed design method might be adapted to other radial filter types. Special treatment, such as regularisation, is then required for radial filters in synthesis problems, that theoretically are unstable at zero or infinite frequency.

References

[1] N. Hahn, F. Schultz, and S. Spors, "Band Limited Impulse Invariance Method," in *30th EUSIPCO*, Belgrade, 2022, pp. 209–213.

[2] N. Hahn, F. Schultz, and S. Spors, "Accurate time-domain simulation of spherical microphone arrays," in *Proc. Forum Acusticum*, Torino, 2023, pp. 599–606.

[3] F. Schultz, N. Hahn, and S. Spors, "Time-Domain Radiation Model of a Cap on a Rigid Sphere," in *Proc. 52nd DAGA*, Dresden, 2026, pp. 1616–1619. [Online]. Available: `https://github.com/spatialaudio/daga2026_dode`

[4] R. M. Aarts and A. J. E. M. Janssen, "Sound radiation from a resilient spherical cap on a rigid sphere," *J. Acoust. Soc. Am.*, vol. 127, no. 4, pp. 2262–2273, 2010. DOI: `10.1121/1.3303978`

[5] E. G. Williams, *Fourier Acoustics*. Academic Press, 1999.

[6] F. W. J. Olver et al., *NIST Handbook of Mathematical Functions*. Cambridge University Press, 2010.

[7] N. Agrawal et al., "Design of digital IIR filter: A research survey," *Appl. Acoust.*, vol. 172, p. 107 669, 2021. DOI: `10.1016/j.apacoust.2020.107669`

[8] B. Bank, "Logarithmic frequency scale parallel filter design with complex and magnitude-only specifications," *IEEE Signal Process. Lett.*, vol. 18, no. 2, pp. 138–141, 2011. DOI: `10.1109/LSP.2010.2093892`

[9] A. Antoniou, *Digital Filters: Analysis, Design, and Applications*, 2nd ed. McGraw-Hill, 1993.

[10] A. Antoniou, *Digital Signal Processing: Signals, Systems, and Filters*. McGraw-Hill, 2006.

[11] A. Antoniou and W.-S. Lu, *Practical Optimization: Algorithms and Engineering Applications*, 2nd ed. Springer, 2021.

[12] M. C. Lang, "Least-squares design of IIR filters with prescribed magnitude and phase responses and a pole radius constraint," *IEEE Trans. Signal Process.*, vol. 48, no. 11, pp. 3109–3121, 2000. DOI: `10.1109/78.875468`

[13] M. Lang, "Algorithms for the Constrained Design of Digital Filters with Arbitrary Magnitude and Phase Responses," Ph.D. dissertation, TU Wien, Vienna, Austria, 1999.

[14] X. Lai, J. Cao, and Z. Lin, "Minimax Magnitude Response Approximation of Pole-radius Constrained IIR Digital Filters," in *Proc. ICASSP*, Brighton, 2019, pp. 5491–5495.

[15] R. C. Nongpiur, D. J. Shpak, and A. Antoniou, "Improved Design Method for Nearly Linear-Phase IIR Filters Using Constrained Optimization," *IEEE Trans. Signal Process.*, vol. 61, no. 4, pp. 895–906, 2013. DOI: `10.1109/TSP.2012.2231678`

[16] D. Guindon, D. J. Shpak, and A. Antoniou, "Design Methodology for Nearly Linear-Phase Recursive Digital Filters by Constrained Optimization," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 57, no. 7, pp. 1719–1731, 2010. DOI: `10.1109/TCSI.2009.2035412`