Received: 13 August 2022 | Revised: 18 September 2022 | Accepted: 29 September 2022

IET Signal Processing

DOI: 10.1049/sil2.12166

The Institution of Engineering and Technology WILEY

ORIGINAL RESEARCH

Minimum phase finite impulse response filter design

Jan C. Olivier[1] | **Etienne Barnard**[2]

[1]School of Engineering, University of Tasmania, Hobart, Australia

[2]Multilingual Speech Technologies (MuST), North-West University, Potchefstroom, South Africa

Correspondence
Jan C. Olivier, School of Engineering, University of Tasmania, Hobart, Australia.
Email: jc.olivier@utas.edu.au

Abstract
The design of minimum phase finite impulse response (FIR) filters is considered. The study demonstrates that the residual errors achieved by current state-of-the-art design methods are nowhere near the smallest error possible on a finite resolution digital computer. This is shown to be due to conceptual errors in the literature pertaining to what constitutes a *factorable* linear phase filter. This study shows that factorisation is possible with a zero residual error (in the absence of machine finite resolution error) if the linear operator or matrix representing the linear phase filter is *positive definite*. Methodology is proposed able to design a minimum phase filter that is optimal—in the sense that the residual error is limited only by the finite precision of the digital computer, with no systematic error. The study presents practical application of the proposed methodology by designing two minimum phase Chebyshev FIR filters. Results are compared to state-of-the-art methods from the literature, and it is shown that the proposed methodology is able to reduce currently achievable residual errors by several orders of magnitude.

KEYWORDS
linear phase filter factorisation, minimum phase finite impulse response filter design, polynomial factorisation

1 | INTRODUCTION

The theory and practice of minimum phase systems and their application have been studied for more than 8 decades [1, 2] and are well understood. The potential advantages offered by a minimum phase system depends on the application [3, 4]. For example, minimum phase finite impulse response (FIR) filters require the least number of taps (lowest possible filter order) given the desired frequency response specifications. Moreover they are unconditionally stable when feedback is deployed [5].

Historically the design of minimum phase FIR filters is performed on the spectral domain (Z domain) [3, 6–8]. In terms of computational requirements, the Hilbert transform and the cepstrum are both based on the discrete Fourier transform[1] and require very long FFT's to obtain good FIR filter performance.

A class of design methods based on the factorisation of a linear phase filter also developed on the Z (spectral) domain. The computation of a minimum phase factor based on computing the roots of a polynomial is well understood [2, 3]. In principle computing the roots of the polynomial provides a solution to the factorisation problem, but is practical only for low order filters. For the factorisation of high order filters, Orchard and Wilson showed in 2003 that spectral factorisation requires the solution of a set of non-linear equations [9][2]—referred to in this paper as the O-W equations. Results presented in ref. [9] showed that the O-W equations readily yield solutions for the high order minimum phase filter based on numerical optimisation. However, the residual errors reported in ref. [9] were not significantly better than the errors that can be achieved based on the Hilbert transform, and thus failed to establish spectral factorisation as a superior minimum phase filter design methodology.

[1]Implemented through the fast Fourier transform (FFT)

[2]In [10] Wilson and Orchard presented a detailed analysis of the factorisation problem.

The next step came in 2017, when Kidambi and Antoniou [11] demonstrated that minimum phase FIR filter design based on the O-W equations and factorisation in fact does provide the best performance currently possible, provided state-of-the-art optimisation is deployed. It was demonstrated in ref. [11] that residual errors can be achieved that are orders of magnitude smaller than those obtained through any other method available in the literature. The proposed methodology and numerical results in ref. [11] have not been improved upon since its publication in 2017, and there is no method available in the literature able to yield a Chebyshev minimum phase filter with a smaller residual error than that reported in ref. [11].

This paper demonstrates that the state-of-the-art residual errors reported in ref. [11] are not at the lowest possible value set by finite machine (computer) resolution. The paper shows that residual errors are at the minimum levels possible on a finite resolution machine, if and only if the linear phase filter that is factorised is *positive definite*—not positive semi-definite as proposed in refs. [11, 12]. The reason for requiring a positive definite linear phase filter becomes clear if the factorisation procedure is analysed in the discrete time domain. The paper derives analytical results based on linear algebra, and shows that a positive semi-definite linear phase filter is not sufficient in general. Numerical results are provided to show that for low order minimum phase filters a positive semi-definite linear phase filter may be able to provide optimal results, but not for high order filters. For the latter case a positive definite linear phase filter provides optimal results, in the sense that residual errors are at a lower limit set by the finite resolution of the digital computer.

This paper will demonstrate that the discrete time domain formulation is able to factorise the difficult case of the Chebyshev approximation based on matrix regularisation. For spectral factorisation the equivalent case was originally considered by Herrmann et al., known as the "lifting" procedure [12]. It will be shown that regularisation implicitly performs lifting, because of the requirement of having a positive definite Hermitian matrix $\mathbf{W}$ representing the linear phase filter $\mathbf{w}$ on the discrete time domain. However, there are subtle differences between the lifting procedure [12] and matrix regularisation, which will be demonstrated in this paper to come down to lifting providing a positive semi-definite linear phase filter. It will also be shown that in the limit where the matrix $\mathbf{W}$ is large (infinite), the O-W equations are in fact solved through Cholesky factorisation of $\mathbf{W}$. Finally, numerical results will be presented to show that a minimum phase FIR design based on a positive definite linear phase FIR filter $\mathbf{w}$, reduces state-of-the-art residual errors reported in ref. [11] significantly. The results provided in this paper establishes the design of a minimum phase filter based on factorisation and the O-W equations as an optimal design methodology, with residual errors many orders of magnitude better than the Hilbert transform.

The paper is structured as follows. In Section 2, discrete time domain factorisation and positive definite minimum phase filter design is proposed, based on locally Toeplitz matrices that are defined. Section 3 presents numerical results for the design of two minimum phase FIR filters, based on 64 bit MATLAB software. The paper is concluded in Section 4.

2 | TIME DOMAIN FACTORISATION AND THE MINIMUM PHASE FILTER c

The passband and stopband specifications for the minimum phase FIR filter $\mathbf{c}$ are the starting point for the design. These specifications are converted to suitable specifications for the linear phase filter $\mathbf{g}$ to be factorised, yielding a minimum phase filter $\mathbf{c}$. The design of an optimal linear phase filter is well understood and was presented in detail in ref. [11], and will not be repeated in this paper. This section analyses the issues associated with the factorisation of $\mathbf{g}$.

2.1 | Locally toeplitz matrices and symmetry point equilibrium

In anticipation that Cholesky decomposition will be required to factorise on the time domain [13], this subsection proposes the deployment of locally Toeplitz matrices that operate on vectors with only local support near the symmetry point.

Consider the transpose of an upper triangular matrix obtained through Cholesky decomposition, as shown in Figure 1. The matrix is not strictly Toeplitz, as the top rows and the bottom rows differ from the rows near the symmetry point, where it is *locally Toeplitz*. This is typical when the matrix size is significantly larger than the number of diagonals that contain non-zero values. Such a matrix is exhibiting symmetry point equilibrium.

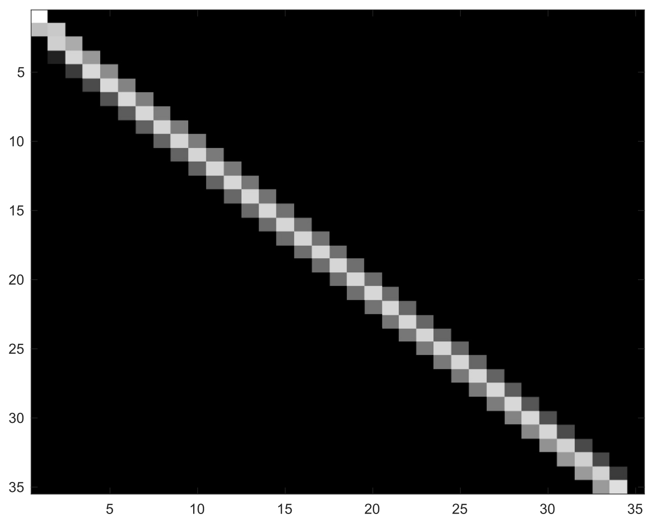

FIGURE 1 The structure of a matrix where symmetry point equilibrium holds—it is locally Toeplitz near the symmetry point, but not strictly Toeplitz.

A matrix that is locally Toeplitz is able to represent a *time-invariant filter*, provided it is operating on a vector that has *only local support* near the symmetry point. To make this statement clear, define an augmented FIR filter so that it has only local support, denoted as $\mathbf{h}_{\text{aug}}$ and defined as[3]

$$\mathbf{h}_{\text{aug}} = \Big\{ \underbrace{0, \dots, 0}_{Q \text{ zeros}}, h_{M-1}, \dots, h_1, h_0, \underbrace{0, \dots, 0}_{Q \text{ zeros}} \Big\}^{\text{T}} \tag{1}$$

where $Q \gg M$. As $\mathbf{h}_{\text{aug}}$ has support only near the symmetry point, a suitable locally Toeplitz matrix where symmetry point equilibrium holds is able to perform a time-invariant convolution operation on $\mathbf{h}_{\text{aug}}$, even if the matrix is not strictly Toeplitz.

To perform convolution, define a Toeplitz matrix with $(2Q + M)$ columns and rows, given by

$$\mathbf{H} = \begin{bmatrix} h_0 & h_1 & \cdots & h_{M-1} & 0 & 0 & \cdots & 0 \\ 0 & h_0 & h_1 & \cdots & h_{M-1} & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & \cdots & 0 & h_0 & h_1 & \cdots & \cdots & 0 \\ 0 & 0 & \cdots & 0 & h_0 & h_1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & \cdots & \cdots & \cdots & \cdots & \cdots & 0 & h_0 \end{bmatrix}. \tag{2}$$

Now define an augmented Kronecker delta as

$$\delta_{\text{aug}} = \Big\{ \underbrace{0, \dots, 0}_{Q \text{ zeros}}, \underbrace{0, \dots, 0}_{M-1 \text{ zeros}}, 1, \underbrace{0, \dots, 0}_{Q \text{ zeros}} \Big\}^{\text{T}} \tag{3}$$

then it follows that

$$\mathbf{h}_{\text{aug}} = \mathbf{H}\ \delta_{\text{aug}}. \tag{4}$$

Define the Gramian matrix $\mathbf{G}$ as

$$\mathbf{G} = \mathbf{H}^{\text{T}}\,\mathbf{H} \tag{5}$$

then it follows that

$$\mathbf{H}^{\text{T}}\ \mathbf{h}_{\text{aug}} = \mathbf{H}^{\text{T}}\,\mathbf{H}\ \delta_{\text{aug}} = \mathbf{G}\ \delta_{\text{aug}}. \tag{6}$$

$\mathbf{G}$ is symmetric, Toeplitz and Hermitian, and represents a linear phase FIR filter [14] denoted as $\mathbf{g}$. The matrix $\mathbf{G}$ has causal non-zero diagonals (the main diagonal and upper triangular part), as well as non-zero diagonals that are anti-causal (the lower triangular part). The dominant (centre) tap of $\mathbf{g}$ is represented by the main diagonal of $\mathbf{G}$.

[3]The fact that the vector has time n advancing from right to left is in anticipation of the fact that Cholesky factorisation implemented by most computing platforms provides a factorisation in the form $\mathbf{U}^T\mathbf{U}$ where $\mathbf{U}$ is upper triangular.

2.2 | Factorisation and regularisation

The matrix $\mathbf{G}$ has real eigenvalues [15], but depending on the application $\mathbf{G}$ may not be positive definite[4]. It is known that in order to be factorable through Cholesky decomposition, the minimum eigenvalue of the matrix must satisfy [15].

$$\lambda_{\min} > 0, \tag{7}$$

that is, it must be positive definite. If Equation (7) is not satisfied by $\mathbf{G}$, then matrix $\mathbf{G}$ must be regularised [15, 16], with the regularised matrix denoted as $\mathbf{W}$ and given by

$$\mathbf{W} = \mathbf{G} + \gamma\mathbf{I}. \tag{8}$$

$\mathbf{W}$ is guaranteed to be factorable as

$$\mathbf{W} = \mathbf{C}^{\text{T}}\mathbf{C} \tag{9}$$

if, and only if γ is chosen as [15, 16].

$$\gamma > |\lambda_{\min}|. \tag{10}$$

The linear phase filter $\mathbf{w}$ is represented by the positive definite matrix $\mathbf{W}$, and can be factorised.

On the spectral domain the process of computing $\mathbf{w}$ is known as lifting [12], but in Section 2.4 it is shown that there is a subtle difference between lifting [12] and regularisation.

2.3 | An approximate minimum phase finite impulse response filter $\mathbf{c}_{\text{approx}}$

The positive definite matrix $\mathbf{W}$ will be factorable based on Cholesky factorisation, and will yield an upper triangular matrix $\mathbf{C}$ as [15].

$$\mathbf{W} = \mathbf{C}^{\text{T}}\mathbf{C}. \tag{11}$$

The matrix $\mathbf{C}$ is locally Toeplitz (if Q is sufficiently large), and represents a time-invariant filter when it operates on vectors with local support as defined above.

Defining a unitary matrix $\mathbf{F}$ as

$$\mathbf{F} = \left(\mathbf{C}^{\text{T}}\right)^{-1}\mathbf{H}^{\text{T}} \tag{12}$$

then for a sufficiently large value of Q it follows that

$$\mathbf{c}_{\text{aug}} = \mathbf{C}\ \delta_{\text{aug}} \approx \mathbf{F}\,\mathbf{h}_{\text{aug}} \tag{13}$$

where $\mathbf{c}_{\text{aug}}$ is given by

[4]This is the case when a Chebyshev linear phase filter is designed.

$$\mathbf{c}_{\text{aug}} = \Big\{\underbrace{0,\ldots,0}_{\text{Q zeros}}, c_{M-1},\ldots,c_1,c_0,\underbrace{0,\ldots,0}_{\text{Q zeros}}\Big\}^{\text{T}}. \tag{14}$$

Even though the relationship between $\mathbf{C}\ \delta_{\text{aug}}$ and $\mathbf{F}\ \mathbf{h}_{\text{aug}}$ is theoretically approximate (due to the regularisation of $\mathbf{G}$), in practice the approximation is excellent.

As Cholesky decomposition expresses a Hermitian positive definite matrix as the product of a minimum phase matrix and its match [15] (regardless of the value of Q) it follows that an *approximate* minimum phase filter can be recovered from $\mathbf{c}_{\text{aug}}$ as

$$\mathbf{c}_{\text{approx}} = \{c_0, c_1, c_2, \ldots, c_{M-1}\}^{\text{T}}. \tag{15}$$

$\mathbf{c}_{\text{approx}}$ is approximate as Q is finite and $\mathbf{W}$ is based on the regularisation of $\mathbf{G}$, but it will be shown in Section 3 that $\mathbf{c}_{\text{approx}}$ is remarkably accurate, even for moderate settings of Q. Section 3 will also demonstrate that $\mathbf{c}_{\text{approx}}$ is an appropriate choice as an initial starting point to perform an efficient numerical optimisation of the non-linear O-W equations to be derived below.

2.4 | The limit where $Q \to \infty$, and the O-W equations

When $Q \to \infty$ symmetry point equilibrium clearly holds, and $\mathbf{C}$ is locally Toeplitz near the symmetry point. Based on the matrix product given by (16), any row near the symmetry point of $\mathbf{C}^{\text{T}}$, say row j, multiplied with columns $j, j+1, j+2, \cdots$ of $\mathbf{C}$ yields a system of non-linear equations given by Equation (17).

$$\mathbf{C}^{\text{T}}\mathbf{C}, \tag{16}$$

$$\begin{aligned} c_0^2 + \cdots + c_{M-1}^2 &= w[M-1] = g[M-1] + \gamma \\ c_0c_1 + \cdots + c_{M-2}c_{M-1} &= w[M-2] = g[M-2] \\ c_0c_2 + \cdots + c_{M-3}c_{M-1} &= w[M-3] = g[M-3] \\ \vdots &= \vdots \\ c_0c_{M-1} &= w[0] = g[0]. \end{aligned} \tag{17}$$

This system of non-linear equations is identical to the O-W equations first proposed in ref. [9] where spectral factorisation was deployed. Hence, if $\lim_{Q\to\infty}$ then Cholesky factorisation solves the O-W equations. This result shows that the O-W equations are in fact exactly solvable (only limited by finite resolution error on the digital machine), if and only if the linear phase filter $\mathbf{w}$ is positive definite. The lifting procedure proposed in refs. [11, 12] requires only a positive semi-definite filter $\mathbf{w}$, for which Cholesky factorisation will fail (as the smallest eigenvalue will be 0). In Section 3 it will be shown that the stricter requirement of a positive definite filter $\mathbf{w}$ leads to a significant reduction in the residual error, which is defined below in Section 2.5. Note that $\mathbf{F}$ is a unitary matrix and an all pass filter[5], hence it follows that as expected.

$$\lim_{Q\to\infty} |C(\Omega)| = |H(\Omega)| \tag{18}$$

TABLE 1 A low order linear phase filter $\mathbf{g}$

Tap number n	$\mathbf{g}$
0	0.066075742625345
1	0.239064282650394
2	0.347182106755652
3	0.239064282650394
4	0.066075742625345

2.5 | Regularisation, lifting, factorisation, and the residual error

Define the residual error vector $\mathbf{e}$ of the O-W equations as

$$\begin{aligned} c_0^2 + c_1^2 + \cdots + c_{M-1}^2 - w[M-1] &= e_0 \\ c_0c_1 + c_1c_2 + \cdots + c_{M-2}c_{M-1} - w[M-2] &= e_1 \\ c_0c_2 + c_1c_3 + \cdots + c_{M-3}c_{M-1} - w[M-3] &= e_2 \\ \vdots &= \vdots \\ c_0c_{M-1} - w[0] &= e_{M-1}. \end{aligned} \tag{19}$$

E_{L_2} denotes the norm of the residual error [11] given by

$$\text{E}_{L_2} = \|\mathbf{e}\|. \tag{20}$$

The norm of any row[6] near the symmetry point of the error matrix given by yields an estimate of the residual error[7].

$$\mathbf{\Delta} = \mathbf{W} - \mathbf{C}^{\text{T}}\mathbf{C} \tag{21}$$

Consider a low order linear phase filter $\mathbf{g}$ shown in Table 1. For the linear phase filter $\mathbf{g}$, the lifting value τ based on the stopband negative ripple peak value (as shown in Figure 2) is given by

$$\tau = 0.00120505352635249. \tag{22}$$

The effect of adjusting the filter $\mathbf{g}$ by adding the lifting value τ to the dominant or centre tap g [2] is also shown in Figure 2. The zeros on the unit circle migrate to form pairs, and hence the adjusted filter (which is positive semi-definite) can be factorised. It will be shown below (and in Section 3) that practical minimum phase filters of high order do *not* factorise

[5] The proof is straightforward and thus not included in this paper.
[6] Using only the upper triangular part and the diagonal.
[7] If Q is chosen sufficiently large the estimate is accurate.

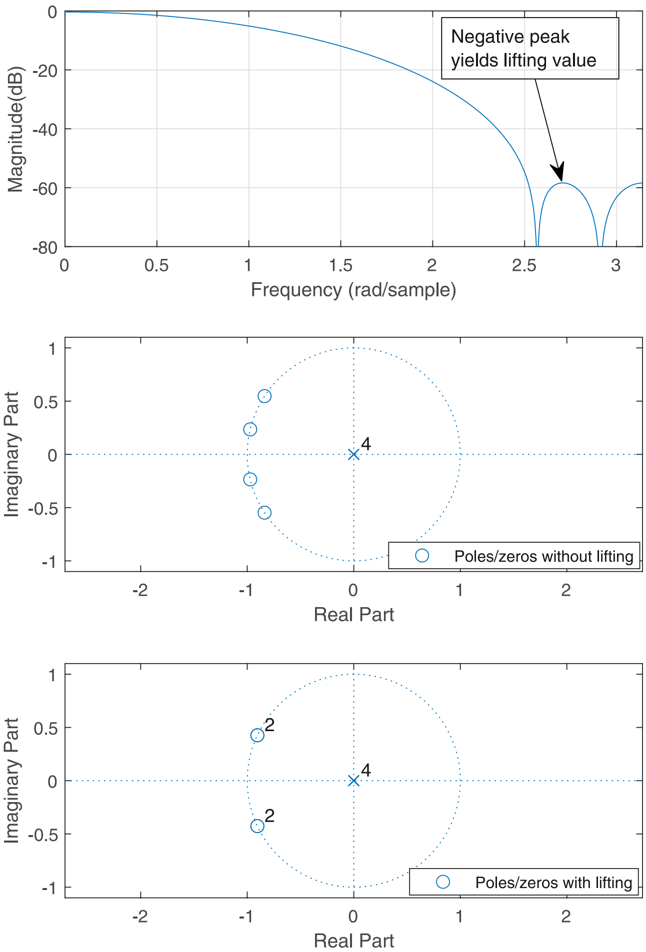

FIGURE 2 The frequency domain and Z plane representation of filter **g** shown in Table 1. The lifting value is provided by the peak value of the negative sidelobe as indicated. In the bottom Figure the lifting value τ = 0.00,120,505,352,635,249 was added to the third tap, which migrated the zeros to form pairs on the unit circle.

if lifting is applied, and that only a positive definite filter **w** can be factorised under those conditions.

The relationship between the lifting value and the minimum eigenvalue is now considered in some detail. The minimum eigenvalue of **G** is shown in Figure 3 as a function of Q, and is negative since the matrix is not positive definite. It is clear that as $Q \rightarrow \infty$ the absolute value of the smallest eigenvalue is converging to the lifting value τ, so that

$$|\lambda_{\min}| = \tau \text{ if } Q \rightarrow \infty. \tag{23}$$

Regularisation requires $\gamma > |\lambda_{\min}| = \tau$ if $Q \rightarrow \infty$, and thus it follows that

$$\gamma = \tau + \epsilon, \tag{24}$$

where ϵ denotes a small positive real number. Unfortunately there is no known method for computing ϵ analytically. Based on numerical studies, it will be shown below and in Section 3 that ϵ is a function of the order of the filter **g**, and must be

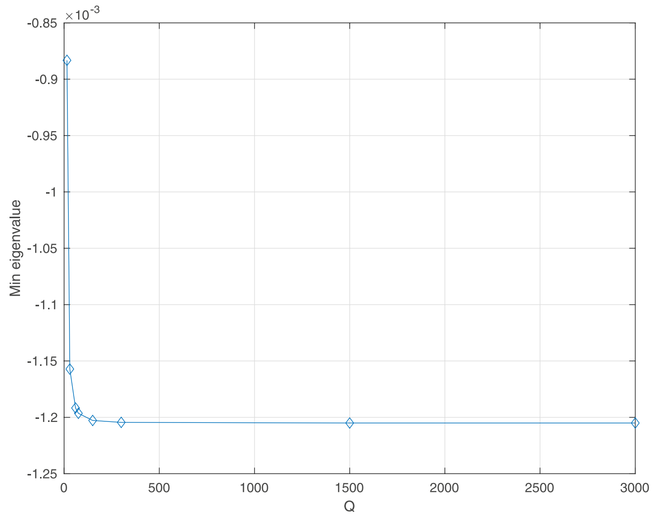

FIGURE 3 The minimum eigenvalue of **G** as a function of Q, for the filter **g** shown in Table 1.

determined for each filter design. If the filter order is small then ϵ is negligible as is the case for the filter shown in Table 1. But it will be shown[8] in Section 3 that $\epsilon > 0$ even for a filter order of 24.

Hence, it follows that the difference between regularisation and lifting comes down to the choice of ϵ. Regularisation requires (in general) $\epsilon > 0$, while lifting [11, 12] deploys $\epsilon = 0$. As is shown in Figure 2, lifting locates the zeros on the unit circle, and thus it is clear that regularisation migrates the zeros (of each pair) away from the unit circle by a very small amount, determined by ϵ. It will be shown in Section 3 that this subtle difference provides a significant reduction in the residual error E_{L_2} for a practical minimum phase FIR filter design.

2.6 | The definition of a factorable linear phase filter w

In this paper, a factorable linear phase filter is defined as follows:

Definition 1 *A linear phase filter* **w** *is factorable, if and only if,* $E_{L_2} = 0$ *and* **c** *is real.*

In practice on a digital computer with finite resolution, the error will of course not be zero, but at an error floor set by the machine finite word length.

It will now be demonstrated (and below in Section 3) that the residual error plays a fundamental role in factorisation. Consider the 5 tap linear phase filter shown in Table 1. Solving the O-W equations based on a 64 bit version of the Levenberg-Marquardt optimiser available in MATLAB as *lsqnonlin* (including the Jacobian) yields a residual error for lifting, that is $\epsilon = 0$, as

[8] In Section 3 a general procedure to determine ϵ will be proposed.

$$E_{L_2} = 3.1032 \times 10^{-17}. \quad (25)$$

This error is at the limit of the smallest error obtainable on a 64 bit version of MATLAB, and thus the lifting value τ is sufficient in this case to guarantee that **w** can be factored.

It will be shown in Section 3 that for higher order filters this is *not* the case, and it is required to deploy a positive definite filter **w** with $\epsilon > 0$ before the residual error is minimised (at the smallest level possible on a finite resolution machine).

3 | CHEBYSHEV MINIMUM PHASE FINITE IMPULSE RESPONSE FILTER DESIGN

The theoretical results presented in Section 2 and new results to be presented in this section can be summarised as follows:

(1) The lifting value τ [12] does *not* yield a factorable linear phase FIR filter for a practical FIR filter order. This section demonstrates this claim, and shows that the state-of-the-art residual errors based on lifting as reported in [11] are not at the lowest values that can be achieved with the 64 bit Levenberg-Marquardt optimiser available in MATLAB.
(2) Based on a discrete time-domain analysis, Section 2 showed that regularisation requires the centre tap of a linear phase FIR filter **g** be increased by a value $\gamma = \tau + \epsilon$, where ϵ is a finite positive (but small) number, providing a positive definite filter **w**. This section proposes to plot the residual error E_{L_2} as a function of γ to determine ϵ—the correct value for ϵ is clearly visible at the waterfall point in the plot, where the error falls away to an error floor.
(3) The residual error E_{L_2} achieved at the waterfall point is the lower limit set by the finite resolution of the digital computer, which for a 64 bit MATLAB Levenberg-Marquardt optimiser is approximately 10^{-17} for the two designs presented in this section.

Two minimum phase FIR designs are presented in this section to confirm the claims above, and results will be compared to current state-of-the-art residual errors reported in ref. [11].

3.1 | A first minimum phase finite impulse response filter design

Example 1 in ref. [11] is deployed in this subsection, and serves to demonstrate the proposed minimum phase FIR filter design procedure. The filter specifications are as follows:

(1) The passband ripple is required to be ≤ 0.173 dB.
(2) The stopband attenuation must be ≤ -50 dB.
(3) The passband cutoff frequency is 0.3π rad/s.
(4) The stopband cutoff frequency is 0.6π rad/s.
(5) The sample frequency is 2π rad/s.

TABLE 2 First half of symmetric tap settings for **g** from ref. [11]

Tap number n	g
0	−0.00033409853951949
1	−0.002489549410806
2	−0.007656350824928
3	−0.011354989160955
4	−0.002981767473881
5	0.018180581093311
6	0.026333770707396
7	−0.008295888670961
8	−0.062043244763120
9	−0.047371546549295
10	0.095349066618093
11	0.295504051520742
12	0.391016383693520

Based on these specifications the tap settings for the linear phase filter **g** is shown in Table 2[9]. The frequency response and the pole/zero representation of **g** are shown in Figure 4. It is clear that **g** cannot be factored—and to confirm this, note that **G** is not positive definite with a minimum eigenvalue that is negative.

To mitigate the non-factorability of the linear phase filter **g**, regularisation is deployed to compute a factorable linear phase filter[10] **w**. The stopband ripple peak measured as $|G(\Omega = \pi)|$ yielded 5.8,174,975 $\times 10^{-6}$. However, some of the negative ripple peaks exceed this value, with the first negative stopband ripple peak most negative, yielding a lifting constant given by

$$\tau = 5.8322404 \times 10^{-6}. \quad (26)$$

As proposed above, the residual error E_{L_2} is plotted over a range of values for γ near τ, as shown in Figure 5. The existence of a waterfall point where the error falls away to a numerical error floor is evident—this is where $\gamma = 5.8{,}322{,}406 \times 10^{-6}$ and thus ϵ can be computed as

$$\epsilon = \gamma - \tau = 5.8322406 \times 10^{-6} - 5.8322404 \times 10^{-6} = 2 \times 10^{-13}. \quad (27)$$

It is clear that $\gamma > \tau$, rendering **w** factorable and positive definite. This empirical result confirms that even for a 25 tap filter the value of γ must *exceed* the lifting value. The initial setting for **c** used in the Levenberg-Marquardt optimiser is

[9]The design procedure to compute **g** for these specifications is presented in detail in [11] and is not repeated here.
[10]Note that additional scaling of **w** is applied to correct for the effects of regularisation on the desired ripple level. The scaling constant is provided in [11] as Equation (8).

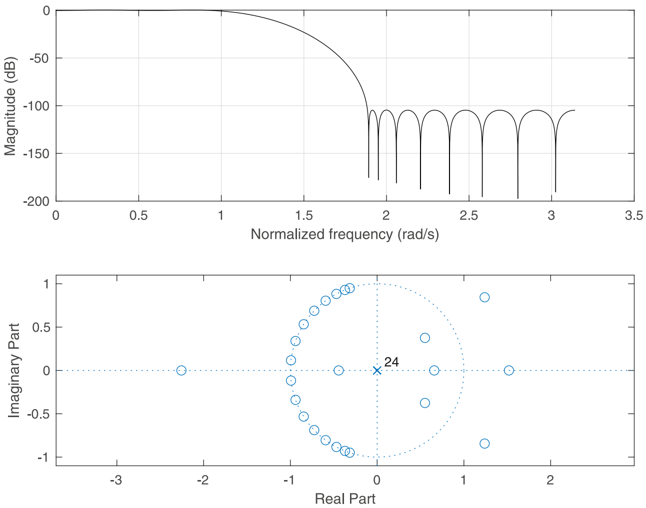

FIGURE 4 A lowpass linear phase filter **g** to be factorised in order to compute the minimum phase finite impulse response (FIR) filter **c**.

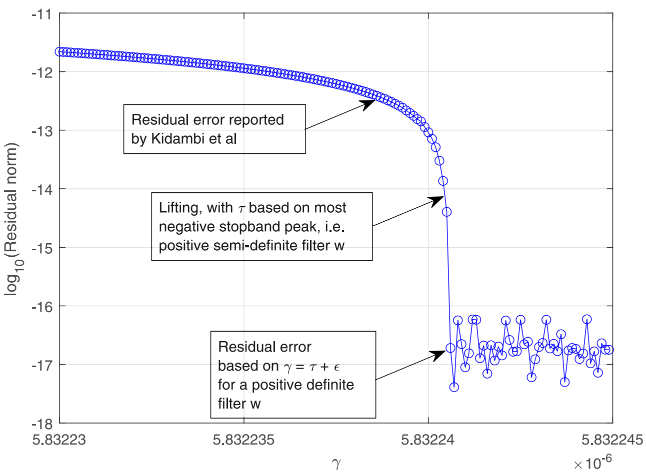

FIGURE 5 The residual error E_{L_2} as a function of γ. Note the waterfall point, which occurs as soon as **w** becomes positive definite as required by regularisation. The residual error reported by Kidambi et al. [11] is also shown.

$\mathbf{c}_{\text{approx}}$, and was computed on the time domain using Cholesky decomposition, with $Q = 10N$ and $N = 25$ for this case[11].

After the waterfall point, the residual error is limited by the 64 bit resolution error floor, and increasing γ beyond this value would simply increase the stop-band leakage, and is thus counterproductive.

Hence, a factorable linear phase filter **w** follows after adding

$$\gamma = 5.8322406 \times 10^{-6} > \tau \tag{28}$$

to the dominant tap of **g**, and yields a linear phase FIR filter **w** that is positive definite and factorable.

[11] N designates the number of taps in **g**

TABLE 3 Settings for the approximate $\mathbf{c}_{\text{approx}}$ and the optimal tap settings for **c**

Tap n	c	$\mathbf{c}_{\text{approx}}$ ($Q = 10\ N$)	c from [11]
0	0.051130001720192	0.053170658603589	0.05111407271
1	0.200742598769625	0.206160489181734	0.2007050616
2	0.373694012448270	0.378228787870903	0.3736680124
3	0.383736569889861	0.380981033331147	0.3837574877
4	0.168053131389221	0.159498759291175	0.1681105336
5	−0.081242863981692	−0.086362886469131	−0.0812083306
6	−0.139775000147583	−0.137267954224048	−0.1397903352
7	−0.028386855628708	−0.024127759326649	−0.0284408539
8	0.060848164109326	0.061160436307306	0.06084337463
9	0.040646929183158	0.038606283413641	0.04066023143
10	−0.011544002965973	−0.012290539783330	−0.01153731949
11	−0.023038330079378	−0.022479612134730	−0.02304231216
12	−0.006534916292685	−0.006283513281458	−0.006536952812

The residual error reported by Kidambi et al. [11] (see Figure 5) is four orders of magnitude above the 64 bit resolution lower limit of $\approx 1.9 \times 10^{-17}$. Note that in spite of this fact, the results reported in ref. [11] are state-of-the-art at the time of writing.

The optimum minimum phase FIR filter **c** obtained by solving the O-W equations with $\gamma = 5.8{,}322{,}406 \times 10^{-6}$ is shown in Table 3. The tap settings for **c** differ in the fifth decimal place when compared to the values reported in ref. [11]. In practice accuracy of five decimal places for the tap settings of a FIR filter is important.

The residual error is shown in Table 4 for a number of methods available in the literature [11], and compared to the residual error based on a positive definite and factorable linear phase FIR filter **w** as proposed in this paper. The method proposed in this paper outperforms the current state-of-the-art method available in ref. [11] by four orders of magnitude for this minimum phase FIR filter design.

The spectral response of the minimum phase filter is shown in Figure 6, indicating that the specifications for the minimum phase filter are met.

3.2 | A second minimum phase finite impulse response filter design

In this section, a second minimum phase FIR design is considered, based on more demanding specifications. In this case the transition region is required to be smaller, and the stopband ripple has been reduced to −70 dB. Thus the specifications for the filter are as follows:

(1) The passband ripple is required to be ≤ 0.173 dB.
(2) The stopband attenuation must be ≤ -70 dB.
(3) The passband cutoff frequency is 0.4π rad/s.

(4) The stopband cutoff frequency is 0.6π rad/s.
(5) The sample frequency is 2π rad/s.

These specifications are met by a 49 tap linear phase filter, with a frequency response shown in Figure 7. The maximum negative ripple level is indicated in Figure 7, and provides a lifting value given by $\tau = 5.106{,}014 \times 10^{-8}$.

Following the procedure proposed in this paper, the value of the residual error E_{L_2} is plotted as a function of γ as shown in Figure 8. The waterfall point occurs when

$$\gamma = \tau + \epsilon = \tau + 9.56015 \times 10^{-12}. \tag{29}$$

TABLE 4 Residual error norm for various methods and the proposed method

Proposed	From [11]	Cepstrum	Hilbert transform	Orchard [9]
1.9×10^{-17}	3.4×10^{-13}	2.4×10^{-6}	2.8×10^{-10}	3.4×10^{-10}

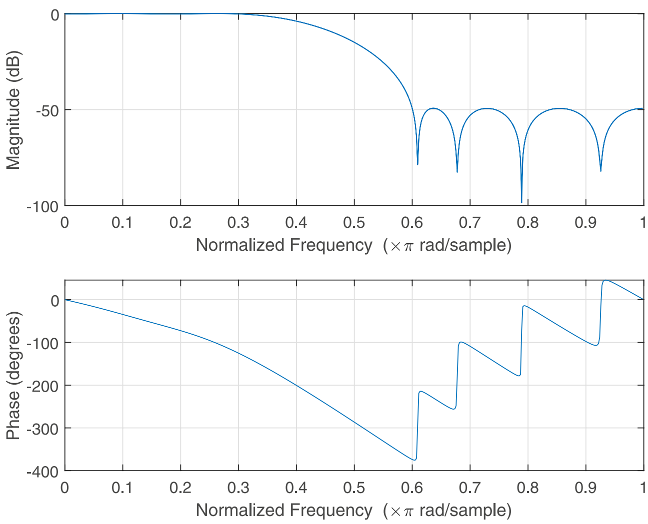

FIGURE 6 The frequency response of the minimum phase filter **c**.

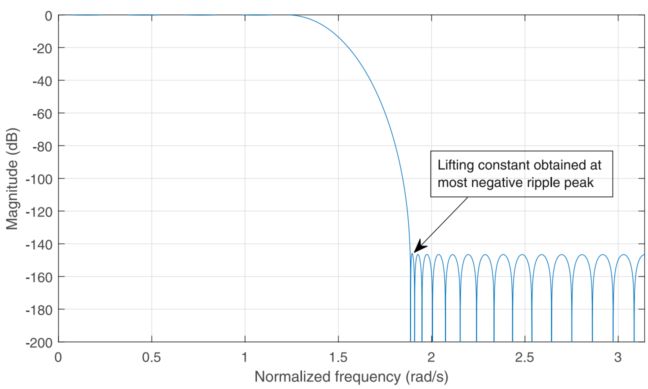

FIGURE 7 The determination of the lifting value for the second example design based on the frequency response of the linear phase filter. The most negative ripple peak occurs at a frequency of $\Omega = 1.895{,}502$ rad/s.

It is clear that $\gamma > \tau$ which renders the linear phase filter factorable, as the value of the residual error E_{L_2} is at the lowest level possible on the 64 bit MATLAB software used to perform the optimisation. The residual error E_{L_2} based on the discrete time domain error matrix $\mathbf{\Delta}$ (see Section 2.5) is given by 1.303×10^{-17}, confirming the residual error provided by the optimisation of the O-W equations.

Note that the lifting value is not able to provide a factorable linear phase filter, as shown in Figure 8. The minimum phase filter **c** follows after Levenberg-Marquard optimisation with γ given above, and the transfer function of the minimum phase filter is shown in Figure 9.

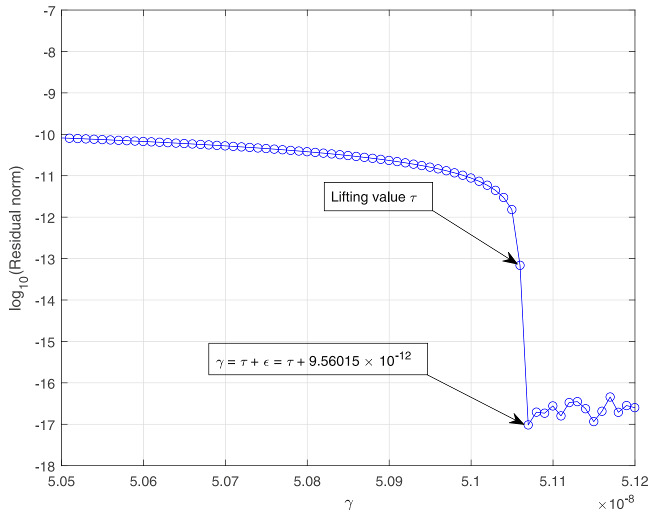

FIGURE 8 The determination of ϵ for the second minimum phase finite impulse response (FIR) filter design, based on the waterfall point proposed in this paper. Note that the lifting constant τ is suboptimal, and does not factorise the linear phase filter **w**.

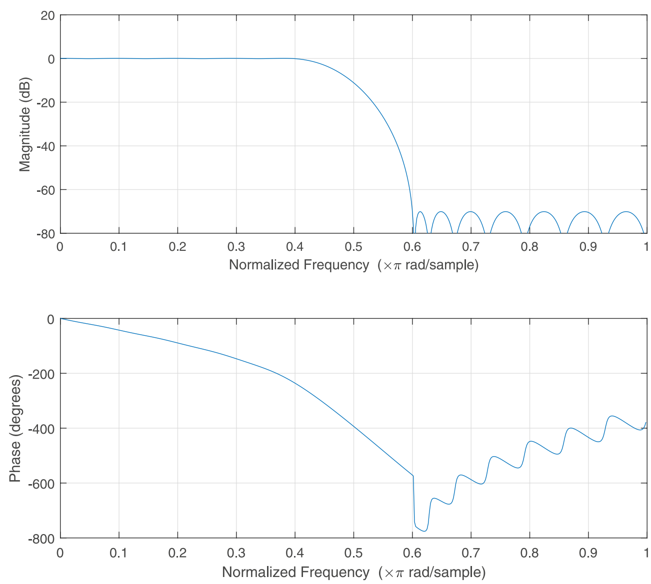

FIGURE 9 The minimum phase filter transfer function, where it is clear that the specifications are satisfied.

3.3 | Complexity analysis of the proposed design procedure

The Levenberg-Marquardt optimisation used to solve the O-W equations is efficient and a low complexity procedure if an accurate initial guess is provided. This paper proposed that the initial guess is based on $\mathbf{c}_{\text{approx}}$, which is based on Cholesky decomposition of the positive definite matrix $\mathbf{W}$. The larger the value of Q, the better the initial guess becomes, but then a larger matrix $\mathbf{W}$ has to be factorised as $\mathbf{W} = \mathbf{C}^{\dagger}\mathbf{C}$. In general, Cholesky factorisation requires $\mathcal{O}(Q^3)$ operations, but it should be recognised that the matrix $\mathbf{W}$ is sparse and Toeplitz—and in fact the larger Q is chosen the more sparse it becomes. For the matrix $\mathbf{W}$ the Cholesky decomposition requires $\mathcal{O}(Q^2)$ operations [17].

The complexity of the proposed method is of the same order of magnitude when compared to the method of Kidambi et al. [11], which is also based on the optimisation of the O-W equations. For the proposed method an additional requirement is that of Cholesky decomposition for $\mathbf{W}$, but the accurate initial guess it provides on the discrete time-domain accelerates the optimisation of the O-W equations significantly[12].

In order to compare the complexity to other methods available in the literature that are not based on the O-W equations, the Hilbert transform is selected as it provides the best overall performance. The Hilbert transform requires a fast Fourier transform (FFT) with more than 10^6 points if a small residual error is required. The Hilbert transform is based on an FFT that has a complexity of $\mathcal{O}(N \log N)$, so that for 10^6 points the complexity of the proposed method is reasonable—given the fact that the residual error is reduced by 7 orders of magnitude compared to the Hilbert transform.

4 | CONCLUSIONS

The paper demonstrated that for high order linear phase filters, the adjustment (or lifting) factor τ based on the negative stoppband ripple peak [12] does not produce a factorable equivalent for a linear-phase Chebyshev filter, and is in practice insufficient for that goal. The paper showed that the definition of what constitutes a factorable linear phase filter has to be based on the O-W [9] equations: the non-linear equations proposed in ref. [9] must be exactly solvable (up to machine precision error) by a minimum phase filter, before the linear phase filter is deemed factorable.

The paper demonstrated that this definition of factorability in practice demands that the linear phase filter $\mathbf{w}$ be positive definite, *not* positive semi-definite as proposed in [11, 12]. Based on a time-domain analysis of factorisation, it was demonstrated that the non-linear equations proposed in ref. [9] are in fact solved through Cholesky factorisation if the linear phase filter is positive definite.

The consequences of a change from a positive semi-definite linear phase filter to a positive definite linear phase filter are dramatic: when the residual error is plotted as a function of the dominant tap adjustment γ, it exhibits a waterfall point coinciding with the point where γ guarantees a positive definite linear phase filter $\mathbf{w}$. At that point the residual error falls away to an error floor set by the finite resolution of the digital computer.

Time-domain factorisation was shown to be possible if an augmentation of the impulse response is defined so that it has only local support, and becomes exact as the expansion approaches infinity. However, even for relatively small expansion factors useful approximations for the minimum phase factor are obtained. This served to provide a good initial starting point for efficient optimisation of the non-linear equations proposed in ref. [9].

Numerical results were provided indicating that the proposed factorisation yields a significant improvement in the residual error when compared to best practice available in the literature.

[12] Kidambi et al. [11] made use of the method proposed in [4] to provide an initial guess.

AUTHOR CONTRIBUTION

Jan C. Olivier and Etienne Barnard contributed to the design and implementation of the research, to the analysis of the results and to the writing of the manuscript.

ACKNOWLEDGEMENT

The authors received no financial support for the research, authorship, and/or publication of this article.

CONFLICT OF INTEREST

The authors have no conflict of interest to declare.

DATA AVAILABILITY STATEMENT

Data available on request from the authors.

ORCID

Jan C. Olivier https://orcid.org/0000-0002-7703-6357

REFERENCES

1. Bode, H.W.: Relations between attenuation and phase in feedback amplifier design. Bell Syst. Techn. J 19(3), 421–454 (1940). https://doi.org/10.1002/j.1538-7305.1940.tb00839.x
2. Oppenheim, A.V., Schafer, R.W., Buck, J.R.: Discrete-Time Signal Processing, 2nd ed. Prentice-hall Englewood Cliffs (1999)
3. Smith, J.O.: Introduction to Digital Filters with Audio Applications. [Online] https://ccrma.stanford.edu/jos/
4. Vainio, O.: Minimum-phase fir filters for delay-constrained noise reduction. IEEE Trans. Instrum. Meas. 48(6), 1100–1102 (1999). https://doi.org/10.1109/19.816120
5. Proakis, J.G., Salehi, M.: Digital Communications. McGraw-Hill, Boston (2008)
6. Venkata, N.D., Evans, B.L., McCaslin, S.R.: Design of optimal minimum-phase digital fir filters using discrete hilbert transforms. IEEE Trans. Signal Process. 48(5), 1491–1495 (2000). https://doi.org/10.1109/78.840000

7. Kamp, Y., Wellekens, C.J.: Optimal design of minimum-phase fir filters. IEEE Trans. Acoust. Speech Signal Process. 31(4), 922–926 (1983). https://doi.org/10.1109/tassp.1983.1164154
8. Reddy, G.R.: Design of minimum-phase fir digital filter through cepstrum. Electron. Lett. 22(23), 1225–1227 (1986). https://doi.org/10.1049/el:19860840
9. Orchard, H.J., Wilson, A.N.: On the computation of a minimum phase spectral factor. IEEE Trans. Circuits Syst. I, Fundam. Theory Appl. 50(3), 365–375 (2003). https://doi.org/10.1109/tcsi.2003.808851
10. Wilson, A.N., Orchard, H.J.: An efficient resultant for determining reciprocal zeros in polynomials. Lin. Algebra Appl. 411, 309–327 (2005). https://doi.org/10.1016/j.laa.2004.08.024
11. Kidambi, S., Antoniou, A.: Design of minimum-phase filters using optimization. IEEE Trans. Circuits. Syst. II: Express Br. 64(4), 472–476 (2017). https://doi.org/10.1109/tcsii.2016.2584002
12. Herrmann, O., Schuessler, H.W.: Design of non-recursive digital filters with minimum phase. Electron. Lett. 6(11), 329–330 (1970). https://doi.org/10.1049/el:19700232
13. Al-Dhahir, N., Cioffi, J.: MMSE decision-feedback equalizers: finite-length results. IEEE Trans. Inf. Theor. 41(4), 961–975 (1995). https://doi.org/10.1109/18.391242
14. Kailath, T.: Linear Systems. Prentice-Hall (1980)
15. Golub, G.H., Van Loan, C.F.: Matrix Computations (2013)
16. https://math.stackexchange.com/questions/3596910/regularization-of-a-matr%20ix-using-a-diagonal-matrix, Accessed 21 09 2021
17. https://stackoverflow.com/questions/69386603/complexity-of-sparse-matrix-%20cholesky-decomposition, Accessed 09 09 2022

How to cite this article: Olivier, J.C., Barnard, E.: Minimum phase finite impulse response filter design. IET Signal Process. e12166 (2022). https://doi.org/10.1049/sil2.12166