Design of Wave Digital Filters with the TU Delft Toolbox

Moritz Tockner*, Hans Georg Brachtendorf†, Martin Steiger‡,
*†‡*University of Applied Sciences Upper Austria*
Softwarepark 11, 4232 Hagenberg, Austria
Corresponding author: *moritz.tockner@students.fh-hagenberg.at

***Abstract*—This paper presents a method for the design of lattice wave digital filters (LWDFs) with an approximately linear phase response. A recently developed algorithm for the design of approximately linear phase allpass filters was utilized in combination with the LWDF structure, to synthesize frequency selective filters with a high phase response flatness. An iterative optimization method to generate such an LWDF for a given set of frequency response constraints was implemented and integrated into an open source LWDF design toolbox to allow an outreach to a broader audience. Additionally, the generated floating point LWDF coefficients can be quantized as multiplier-free signed digits by solving a nonlinear optimization problem. The design toolbox finds the coefficient representation requiring the minimum number of additions and/or subtractions, while still satisfying a given set of frequency response constraints. The results show the effectiveness of the filter design method compared to standard designs such as Butterworth, Tschebyscheff or Cauer and they highlight the differences to another existing linear phase method of the toolbox. A significant improvement can be seen in a comparison between the signed digit quantization compared to simple coefficient rounding.**

***Index Terms*—design method, linear phase filters, lattice wave digital filter, LWDF, signed digit quantization**

I. Introduction

The lattice wave digital filter (LWDF) is a specific implementation structure of an infinite impulse response (IIR) filter. It poses numerous advantages for implementation in hardware such as low coefficient wordlength requirements, low roundoff noise level and no parasitic oscillations, while being highly modular. These properties make the LWDF a favorable choice for very large scale integration (VLSI) [1]. Standard IIR filters exhibit a nonlinear phase response, which makes them unsuited for applications which rely on the shape of the signal envelope. LWDFs however, can be designed to reduce this nonlinearity to a reasonable amount. A new approach to the design problem of reducing this nonlinearity was proposed in [2]. It utilizes an allpass filter design algorithm that maximizes the filters phase response flatness in a specified band, while obtaining an equiripple phase response in the remaining band(s) by applying the Remez exchange algorithm.

The intention of this work is to make this design method conveniently accessible for practical use, by implementing the algorithm into an open source design toolbox that lets the user generate an LWDF with standard design methods like Butterworth, Tschebyscheff or Cauer, for a given set of filter constraints [3]. This toolbox was already extended by another method for achieving approximately constant group delay [4], which will be used for a comparison with the new implementation. Furthermore, the quantization behavior of the resulting filters will be analyzed for a simple method i.e. coefficient rounding, and a more sophisticated method which uses a signed digit representation for the LWDF coefficients. An algorithm for finding the optimum representation with respect to the number of necessary additions and/or subtractions in hardware, for a given set of frequency response constraints, is explained in [5] and will also be made publicly available through the toolbox.

The general principle of the LWDF structure is to create a sum from the output of two allpass filters which ideally exhibit a phase relation of 0 in the passband and $\pm\pi$ in the stopband. A division by 2 is done at the end to ensure that the resulting magnitude response is ≤ 1. The resulting system function can be described as

$$H(z) = \frac{1}{2}\left[A_0(z) + A_1(z)\right], \tag{1}$$

where $A_0(z)$ and $A_1(z)$ are the allpass filter system functions. To achieve a low roundoff noise level, the individual allpass filters are further divided into a cascaded structure consisting of an interconnection of two port adaptors and delay elements. These can be written as

$$H_{1,2}(z) = \prod_{i=1}^{M} H^{(i)}(z) \tag{2}$$

where $H^{(i)}(z)$ is either a first order section

$$H^{(i)}(z) = \frac{-\gamma_i + z^{-1}}{1 - \gamma_i \cdot z^{-1}} \tag{3}$$

which is fully defined by the coefficient γ_i, or a second order section

$$H^{(i)}(z) = \frac{-\gamma_{i0} + \gamma_{i1}(\gamma_{i0} - 1)z^{-1} + z^{-2}}{1 + \gamma_{i1}(\gamma_{i0} - 1)z^{-1} - \gamma_{i0}z^{-2}}, \tag{4}$$

defined by γ_{i0} and γ_{i1}. It can be noted that the linear phase design which will be addressed in the following section uses only delay elements in the lower branch. An example filter with this kind of structure is illustrated in figure 1.

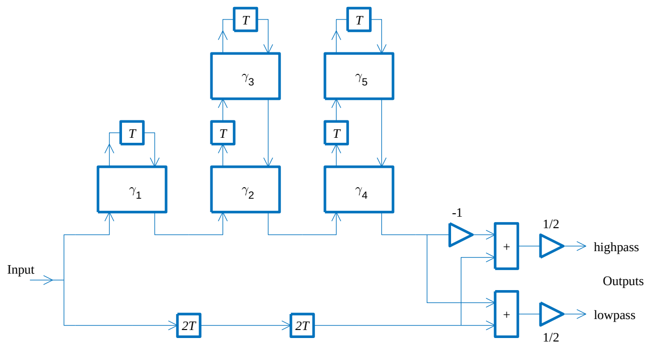

Fig. 1. A toolbox visualization example of a fifth order lowpass LWDF structure.

II. Linear Phase LWDF Design

The design method featured in this paper relies on the work of Zhang [2], [6], [7]. It requires putting initial constraints on the filter's phase response degree of flatness K at one or multiple frequency points ω_p, while using the remaining degrees of freedom for an equiripple phase response in specified band(s). Previously existing methods did not permit the design of allpass filters with these two properties combined. For an allpass filter with the general transfer function

$$A(z) = z^{-N} \frac{\sum_{n=0}^{N} a_n z^n}{\sum_{n=0}^{N} a_n z^{-n}}, \tag{5}$$

the definition for the degree of flatness K ($\in \mathbb{Z}$) at a frequency point ω_p can be expressed as

$$\left.\frac{\partial^r \Theta(\omega)}{\partial \omega^r}\right|_{\omega=\omega_p} = \left.\frac{\partial^r \Theta_d(\omega)}{\partial \omega^r}\right|_{\omega=\omega_p}, \quad r = 0, 1, ..., K-1 \tag{6}$$

where Θ_d and Θ are the desired and actual phase responses. By substitution of the general phase response of an allpass filter, and assuming a linear desired phase response, these conditions can also be stated as

$$\begin{cases} \sum_{n=0}^{N} (n - \frac{N-\tau}{2})^r \sin\{(n - \frac{N-\tau}{2})\omega_p\} a_n = 0 & \text{for even } r, \\ \sum_{n=0}^{N} (n - \frac{N-\tau}{2})^r \cos\{(n - \frac{N-\tau}{2})\omega_p\} a_n = 0 & \text{for odd } r. \end{cases} \tag{7}$$

It can be seen that the equation is always true for odd r if $\omega_p \in \{0, \pi\}$. It follows that that the number of conditions is

$$L = \begin{cases} \lfloor \frac{K}{2} \rfloor & \text{for } w_p \in \{0, \pi\}, \\ K & \text{otherwise.} \end{cases} \tag{8}$$

To generate meaningful results, L needs to be less than the filter order N. Therefore the remaining degree of freedom $N - L + 1$ can be used for the equiripple phase response in the specified band(s) by applying the Remez exchange algorithm. The filter coefficients a_n can then be obtained by solving a generalized eigenvalue problem in the form of

$$\boldsymbol{PA} = \delta \boldsymbol{QA}, \tag{9}$$

where $\boldsymbol{A} = [a_0, a_1, ..., a_N]^T$, and the elements of the matrices $\boldsymbol{P}$ and $\boldsymbol{Q}$ are given by

$$P_{ij} = \begin{cases} \left.\frac{\partial^i \sin\{(n - \frac{N}{2})\omega - \frac{\Theta_d(\omega)}{2}\}}{\partial \omega^i}\right|_{\omega=\omega_p} \\ \qquad \text{for } i = 0, 1, ..., L-1, \\ \sin\{(j - \frac{N}{2}\omega_{(i-L)} - \frac{\Theta_d(\omega_{(i-L)})}{2}\} \\ \qquad \text{for } i = L, L+1, ..., N, \end{cases}$$

$$Q_{ij} = \begin{cases} 0 & \text{for } i = 0, 1, ..., L-1, \\ (-1)^{i-L} \cos\{(j - \frac{N}{2}\omega_{(i-L)} - \frac{\Theta_d(\omega_{(i-L)})}{2}\} \\ & \text{for } i = L, L+1, ..., N. \end{cases} \tag{10}$$

A more detailed derivation can be found in the original paper by Zhang [2].

By using this design method in combination with an LWDF structure, frequency selective filters can be generated which exhibit maximal flatness in their passband(s) and equiripple in their stopband(s) phase response. To obtain a filter that satisfies given design constraints in the form of passband and stopband magnitude response tolerances, an iterative procedure can be used. At first, a coefficient set is generated with the lowest possible order ($N = 2$) and the highest corresponding degree of flatness

$$K = \begin{cases} 2N - 1, & \text{for } 2N - 1 \leq 15, \\ 15 & \text{otherwise.} \end{cases} \tag{11}$$

$K \leq 15$ was found to be a reasonable upper limit during the implementation of the algorithm. If the resulting filter does not satisfy the requirements, K is decremented and the design method is tried again. If the requirements are still not satisfied at the lowest reasonable value $K = 3$, N can be incremented and again all possible K are tested beginning with the highest. This process is repeated until a solution is found or the reasonable upper limit for N is reached. This upper limit was empirically determined to be $N = 20$. For $N > 20$, the algorithm generates filters with poles in the passband, which makes it impossible to fulfill a given ripple tolerance.

III. Signed Digit Quantization

A resource efficient way of implementing this LWDF design in hardware can be done by following the work of Kaakinen and Saramäki [5]. They propose implementing the multiplication via a sequence of shift and add and/or subtract operations, instead of using costly general multipliers. For this purpose, it is desireable to express the coefficient values in the form of

$$\sum_{r=1}^{R} x_r 2^{-P_r}, \tag{12}$$

where $x_r \in \{-1, 0, 1\}$ and $P_r \in \mathbb{N}$.

Here, the difference to the classical canonical signed digit (CSD) representation is the missing property that adjacent CSD digits are never both non-zero. This constraint leads to a bijective conversion between a 2's complement number and the corresponding CSD representation. Therefore, the signed digit approach leads to a non-injective conversion from 2's complement to signed digits, but potentially decreases the necessary number of non-zero digits.

Thus the quantization goal is to find the optimum signed digit representation for a given set of floating point coefficients, with respect to the implementation cost. This means that R needs to be as small as possible. Then, a solution needs to be found for this minimum R with the smallest maximum P_r.

This optimization can be done in two steps. First, for every allpass denominator coefficient, the minimum and maximum possible values that still satisfy the constraints, need to be found while reoptimizing the others. The constraints are composed of three parts. First, magnitude response passband and stopband tolerances

$$\begin{aligned} 1-\delta_p \le |H(e^{j\omega})| \le 1 \quad & \text{for } \omega \in \Omega_p \\ |H(e^{j\omega})| \le \delta_s \quad & \text{for } \omega \in \Omega_s, \end{aligned} \tag{13}$$

where Ω_p and Ω_s are passband and stopband regions and δ_p and δ_s the respective tolerances. Second, phase response linearity tolerance

$$|\arg H(e^{j\omega}) - \tau\omega| \le \Delta \quad \text{for } \omega \in \Omega_p, \tag{14}$$

where τ is the ideal group delay. The last constraint is keeping the order of the filters poles after their radii intact. This problem can be solved by using any nonlinear problem solver.

The next step is then to find the optimum signed digit representation inside these boundaries, according to the aforementioned criteria. This can be done in an exhaustive manner, by testing all possible combinations beginning with the least cost. $R = 1$ and maximum $P_r = 1$, and incrementing first the maximum P_r and then R if the constraints are not satisfied. This approach can however, lead to an unfeasible runtime for some filter constraints.

IV. LWDF Design Toolbox

The Delft University of Technology provides an open source toolbox named Wave Digital Filter Designer [3], which focuses solely on the design of WDFs and LWDFs. It was extended by the work of Zeintl et al. [4] with a different approach for the design of approximately linear phase LWDFs. This made it an ideal candidate for integration of the practical part of this work, because it ensured an easy way of comparing the results, with existing methods.

Figure 2 shows the general design flow of the toolbox, where the proposed changes are drawn with dashed lines. It allows generation of the four general filter types lowpass, highpass, bandpass and stopband. The initial version from TU Delft includes the standard approximation methods Butterworth, Chebyshev, Cauer and the Sharpe/Vlach method

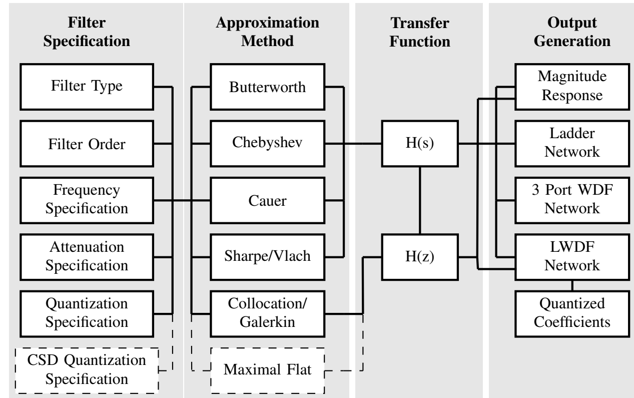

Fig. 2. Current design flow of TU Delfts Toolbox with additions from Zeintl [4] (solid) and proposed changes (dashed).

for a more direct possibility to design passband and stopband filters. The Collocation/Galerkin methods are the linear phase approximations that where added in [4]. The filters can be designed as time-continuous or time-discrete, where the standard methods are transformed from the continuous time domain via the bilinear transformation. The resulting filters' frequency response can be displayed, as well as the actual WDF parameters for two and three port adaptor networks. The parameters of the two port adaptor network can also be quantized with the standard methods like rounding and truncation. For a time-continuous design, a ladder network consisting of resistors, capacitors and inductors can also be created.

V. Results

As a result of this work, the approximately linear phase LWDF design method and the signed digit quantization algorithm were both integrated into the toolbox from [3] to be publicly available. A comparison between a standard Butterworth design, the Collocation method of [4] and the newly implemented maximal flat method can be seen in figure 3 and 4. The filter order was chosen such that their hardware realization as LWDF would lead to $M = 9$ multipliers. For the maximal flat design algorithm, passband and stopband where defined as $\Omega_p = [0, 0.25]$ and $\Omega_s = [0.35, 0.5]$ with magnitude response tolerances $\delta_p = 0.3$ dB and $\delta_s = 30$ dB. The cutoff frequencies for the other two algorithms where chosen to satisfy the same criteria. It can be seen in figure 3 that the Collocation and maximal flat methods produce approximately the same magnitude response, with the Collocation method displaying a slight passband ripple. The maximal flat method on the other hand shows a faster roll-off near the end of the passband, while the Butterworth method exhibits the steepest roll-off of all three methods. The group delay plot in figure 4a shows the impracticality of the Butterworth method, when low group delay variation is needed. The main difference between the Collocation and the maximal flat filters can be seen in figure 4b, where the group delay plot is magnified in the passband. Here, the Collocation design still shows a small

TABLE I
TABLE OF DIFFERENT QUANTIZATIONS FOR LWDF γ COEFFICIENTS.

	Floating-Point (Decimal)	Rounded (Binary string)	Signed Digits (Trinary string)
γ_1	-0.078139174739667	1.111	.00+0
γ_2	0.494223541933836	0.100	.0+0+
γ_3	-0.119310280159436	1.111	.0000
γ_4	0.357791830705940	0.011	.00++
γ_5	-0.216078536280463	1.110	.00-0
γ_6	-0.688360259211653	1.010	.-000
γ_7	-0.629091302981285	1.011	.--00
γ_8	-0.124584853128215	1.111	.0000
	Adders/Subtractors	23	9

ripple, whereas the maximal flat design is perfectly flat. The cost of this flatness can be seen near the end of the passband, where the group delay of the maximal flat design starts to rise earlier than that of the Collocation design.

The signed digit quantization algorithm was compared with a standard coefficient rounding method for a maximal flat design. The filter was designed for $\Omega_p = [0, 0.25]$ and $\Omega_s = [0.35, 0.5]$, with $\delta_p = 1$ dB and $\delta_s = 50$ dB. The quantization was then optimized for $\Omega_{p,q} = [0, 0.20]$ and $\Omega_{s,q} = \Omega_s$, with $\delta_{p,q} = \delta_p$, $\delta_{s,q} = 40$ dB, and $\Delta\tau_q = 0.5$. These constraints are satisfied for a representation with a maximum $P_r = 4$ and a maximum $R = 2$ for all coefficients. Table I shows the floating point coefficients generated by the design algorithm, the optimized signed digit representation and the rounded two's complement representation to a comparable number of bits. The last line of the table shows the number of necessary adders and subtractors that would be needed for an implementation in hardware via shifts and adds or subtracts. Only the adders and subtractors needed for multiplication with the γ coefficents are counted for this comparison. For a realization of the filters, additional structural adders would also be needed. These are for both coefficient representations the same and therefore not included in the comparison. In this regard, the signed digit quantization shows a significant improvement in comparison to the rounding. It also has to be noted that the magnitude response of the signed digit representation behaves drastically better over the whole bandwidth than the rounded version. In this particular case even almost as good as the floating point version, while only adding the previously defined $\Delta\tau_q \leq 0.5$ to the group delay in most of the original passband.

VI. CONCLUSION

In this paper, a recently developed method for the design of approximately linear phase response LWDFs was implemented and integrated into an open source toolbox. The implementation shows an improvement of group delay flatness in the filters passband, when compared to designs of the existing Collocation and Butterworth methods with equal realization effort. Furthermore, a quantization method that minimizes necessary hardware resources by utilizing the signed digit representation was implemented. The filters stopband attenuation and group delay variation degrade less when using this method, than from a simple rounding of its coefficients. This stems from the use of optimization constraints for these properties. Due to the nature of the signed digit optimization, the necessary hardware resources for a realization of the multiplications in the filter structure as shift and add or subtract operation are also typically lower compared to the rounding to a two's complement of similar bit width. The downside of this implementation is a typically high runtime. Depending on the chosen constraints, this can become unpractical due to the exhaustive part of the optimization. Further studies could be directed to the search for a heuristic approach to speed up this process.

ACKNOWLEDGEMENT

This project AMOR ATCZ203 has been co-financed by the European Union using financial means of the European Regional Development Fund (INTERREG) for sustainable cross boarder cooperation. Further information on INTERREG Austria-Czech Republic is available at https://www.at-cz.eu/at.

REFERENCES

[1] L. Gazsi, "Explicit formulas for lattice wave digital filters," *IEEE Transactions on Circuits and Systems*, vol. 32, no. 1, pp. 68–88, 1985.
[2] X. Zhang, "Design of allpass filters with specified degrees of flatness and equiripple phase responses," in *2010 International Conference on Signal Processing and Multimedia Applications (SIGMAP)*, pp. 205–210, 2010.
[3] H. J. Lincklaen Arriëns, "(l)wdf toolbox," 2006.
[4] C. Zeintl and H. Brachtendorf, "Linear phase design of lattice wave digital filters," pp. 1–5, 04 2018.
[5] J. Yli-Kaakinen and T. Saramäki, *A Systematic Algorithm for the Synthesis of Multiplierless Lattice Wave Digital Filters*. IntechOpen, 04 2011.
[6] Xi Zhang and H. Iwakura, "Design of iir digital allpass filters based on eigenvalue problem," *IEEE Transactions on Signal Processing*, vol. 47, no. 2, pp. 554–559, 1999.
[7] Xi Zhang and H. Iwakura, "Design of iir digital filters based on eigenvalue problem," *IEEE Transactions on Signal Processing*, vol. 44, no. 6, pp. 1325–1333, 1996.

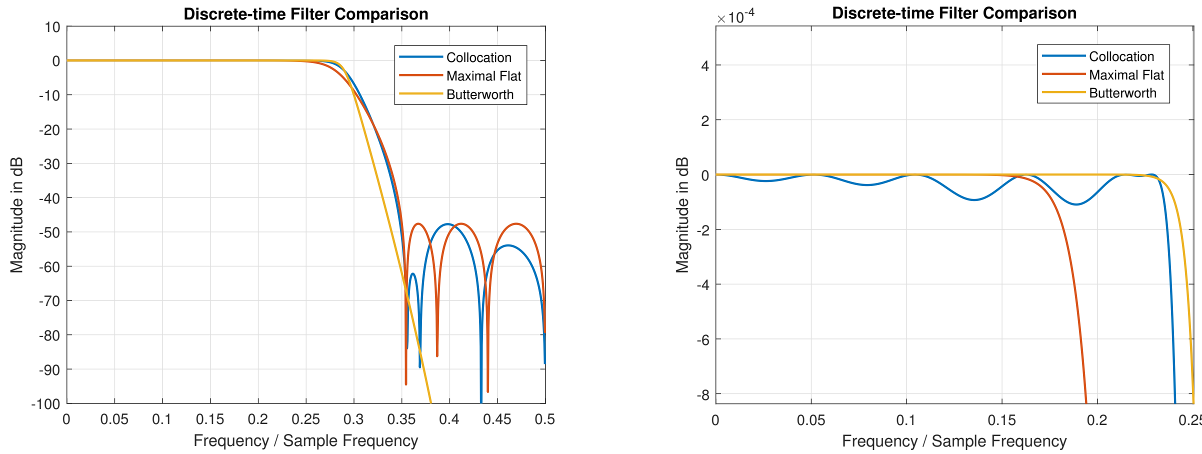

(a) (b)

Fig. 3. Comparison of magnitude responses for lowpass Butterworth, Collocation and maximal flat designs with equal implementation effort (multiplier count $M = 9$) full spectrum (a) and zoomed in on the passband (b).

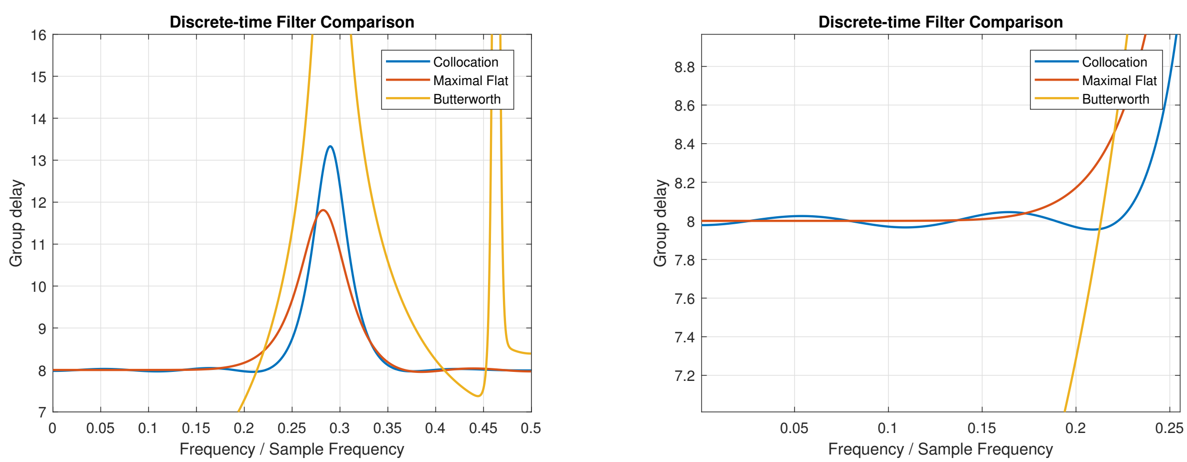

(a) (b)

Fig. 4. Comparison of group delays zoomed in on the passband for lowpass Butterworth, Collocation and maximal flat designs with equal implementation effort (multiplier count $M = 9$) full spectrum (a) and zoomed in on the passband (b).