EQ Structures Comparison

Aleksi Peussa
Aalto Universtiy
Master's Programme CCIS / AAT

aleksi.peussa@aalto.fi

Abstract

This paper reviews the development of graphic equalisers using digital signal processing (DSP) for its two main architectures: parallel and cascade. After over viewing each system, modern developments are investigated which show that both implementations have been able to achieve a sub ±1dB error therefore classifying them as hi-fidelity equalisers. The most significant difference between the two approaches then lies in computational cost, in which improvements are still being made. Lastly, a possible hybridization of the graphic equaliser was investigated using a classic example, but it did not achieve improvements to the modern methods. Additional testing and adjustment of the modern methods to suit the hybrid structure would be necessary. Future implementations of a possible hybrid structure could aim to provide benefits to unresolved issues with graphic equalisers by focusing on reducing computational costs. This could be accomplished by providing the simultaneous computing benefits of the parallel structure but reducing the operations of an entirely parallel system by using cascaded sections.

1 Introduction

Equalisers originated in telephone engineering where the spectrum of the sent signal was adjusted to make up for the transmission losses [1]. This aimed to provide a flat spectrum hence the origin of the term was directly related to producing a flat frequency spectrum. While further developments of this technology did to some extent aim to even out the response one of the earliest cases of how we see equalisers nowadays was the development of the tone control, which was first used on gramophones in the 1940's by means of switches [2]. The next progression of this idea allowed the user to control the filters more by using potentiometers, the objective was to control the frequency response of the playback system by means of two shelving filters; one for the high frequencies and one for the low, i.e. high- and low-shelving filters [3].

While these kinds of tone controls became commonplace in music playback devices, one of the most important developments came when Massenburg (alongside others [1]) published his paper on the parametric equaliser [4]. The parametric equaliser was unique because it occupied a narrower region of the frequency spectrum since it only affected a specific band, which made it different from the previously mentioned shelving filters. However, this was not the only important part of the filter. Its control was simple, having the user decide not only the centre frequency, but also the Q and amplitude/gain of the filter. Massenburg developed a three-band equaliser from this advancement [1], but this idea would be taken far further.

The parametric equaliser would be the fundamental building block for the graphic equaliser, a collection of many fixed Q filters that would be combined to provide an equalisation control of the entire audible frequency range by separating the filters logarithmically to match the nature of human hearing. By controlling each filter with a gain slider, the adjustments to the frequency spectrum could be visualised [1] and these analogue devices became standard in music mixing studios to finalise the production of recordings. While in many of these cases the person mastering will most likely have listened to the music more so than looking at the equaliser levels, it was still expected that the changes graphically would correspond to what was occurring sonically. However, the consecutive filters are inherently not perfectly isolated from each other and result in a spectrum that is far from what the visual interpretation would lead the user to believe [5], or even what they would have wanted.

With the development of digital signal processing (DSP) techniques in the field of equalisation many studies have tried to investigate and delve deeper into graphic equalisers. While in many traditional cases the user may not care about the absolute accuracy of the end spectrum of the equaliser, it still is an engineering goal to make graphic equalisers function so that the desired spectrum is convincingly output by the system. The research into providing the best possible graphic equaliser in both efficiency and accuracy can be divided into two main structures: cascade and parallel. Cascade was the traditional method described earlier, in which fixed Q parametric filters are connected in series to cover the entire audible spectrum [1]. On the other hand, parallel structures divide the input signal to branches that are then bandpass filtered and summed together to provide the output [1].

This paper will focus on comparing these two fundamental structures for designing a digital graphic equaliser by reviewing the current leading research into both design methods. While some conclusions will be made by reviewing literature, a further inspection of possible hybrid structures that combine both parallel and cascade filtering to build the entire graphic equaliser will be performed. This will lead to a current overview of graphic equaliser performance using DSP, while also pursuing the possibility of combining the two approaches to improve overall performance.

2 Equaliser Structures

2.1 Cascaded Graphic Equalisers

As stated prior, the cascaded structure involves chaining multiple fixed Q parametric filters to cover the entire audio spectrum and form the equaliser. Usually the spacing of the filters is logarithmic in order to match them to human hearing frequency resolution, for example an octave spaced graphic equaliser would have 10 bands and a third octave equaliser would have 30 bands (specified in ISO standard [6]). The frequency spacing of these filters can be seen in table 1 below.

Table 1: Octave band and third octave band filter centre frequencies [1]

Octave f_c (Hz)	**$\frac{1}{3}$-Octave f_c (Hz)**	**Octave f_c (Hz)**	**$\frac{1}{3}$-Octave f_c (Hz)**
-	25	-	800
31.5	31.5	1000	1000
-	40	-	1250
-	50	-	1600
63	63	2000	2000
-	80	-	2500
-	100	-	3150
125	125	4000	4000
-	160	-	5000
-	200	-	6300
250	250	8000	8000
-	315	-	10000
-	400	-	12500
500	500	16000	16000
-	630	-	20000

Whichever spacing is chosen for the graphic equaliser the fundamental structure will remain the same with each filter adjusting its respective region with either a peak or notch of varying degree, specified by the user. While each individual filter is specified a command gain, an optional pre-gain may be applied which helps in the condition that all the command gains are the same [1]. It is important to note that if the entire frequency range is covered with only peaking/notching filters there will be 0dB gain at DC and Nyquist [7], which is usually an undesired function. This can be rectified by replacing the first and last parametric filters with low/high shelving filters respectively, which can be either controlled with an individual gain or by setting the shelving gain equal to the neighbouring parametric filter [7]. One benefit of the cascaded structure is that the output will have the properties, poles and zeros, of the individual filters. This means that if the individual filters are all minimum-phase, a desirable trait in audio DSP, then the total output will also be minimum-phase [1]. The generalised structure of the cascade graphic equaliser can be seen below in Fig.1.

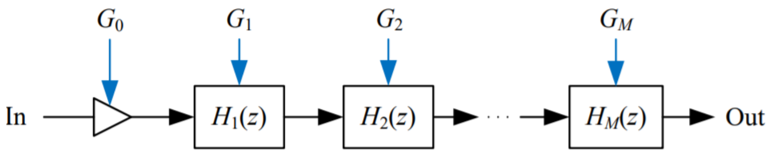

Figure 1: *Cascaded graphic equaliser structure [1]*

Ideally each filter would crossover to the next one without interfering with its individual response, but this is not the case. The decay of each filter will have a significant effect on the overall output, which if not accounted for leads to an equaliser that does not produce the user's command/input gains accurately. Simple adjustments of the individual filters by for example reducing or increasing the Q will not save the performance but will only bring forth its own set of problems. For the prior adjustment the output gain will have dramatic troughs in the overall response between the command gains and for the latter the output gain will be dramatically higher between the command points than intended [7].

2.2 Parallel Graphic Equalisers

In the parallel architecture the individual filters are now bandpass filters that fundamentally produce the desired response by forming a resonance in their pass-band, while having a minimal gain elsewhere [5]. Due to the filters now being in parallel with each other when calculating the response of each filter, both the phase and magnitude response will affect the output spectrum. This leads to a more complicated design process even though the mathematics behind both cascade and parallel should make them identical [1]. The structure of the parallel architecture can be seen below in Fig.2, note also the second order filter structure H_k shown beside it for a fixed pole design [5][8].

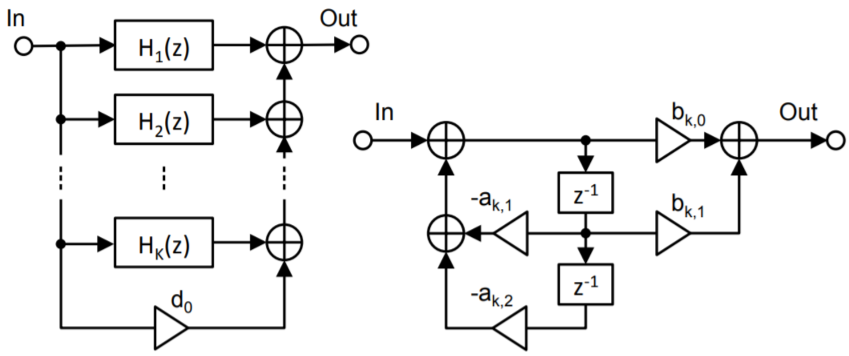

Figure 2: *Parallel graphic equaliser and individual second order filter structure [5]*

One benefit of this structure is that the output is comprised of a single stage of filtering, therefore only one stage of quantisation noise is accumulated unlike the chain of noise formed in the cascaded structure [1]. Furthermore, achieving a flat response across all the command gains is simple with the parallel structure since a pass-through path can be added that only adjusts the gain of the input signal [1]. Designing and optimising parallel graphic equalisers is a more challenging task than its counterpart but can offer very accurate results if accurately designed. Research has also gone into making the parallel structures more optimised and efficient for use in real-time applications over their counterpart structure. One example of this is the use of parallel computing in graphics processing units (GPUs), the parallel structure can take advantage of that by using parallel form IIR filters [9] leading to improved performance.

While parallel computing provides faster performance, a key improvement was the development of a fixed-pole design developed by Bank in 2007 for instrument body modelling [8]. It is practical for other audio filter tasks like equalisation because it provides efficient filtering with a logarithmic spacing of frequencies once the poles have been defined [5][10], note that the general structure for the individual filters was shown in Fig.2 above. Another alternative approach was proposed by Virgulti et al. [11] using higher order IIR filters with downsampling, but these increase the order of the filters from the aforementioned reduced second order structure [1].

3 Current High Performance

Higher performance cascade structures have developed alongside the progression in DSP methods, while also becoming more sophisticated due to improved computation capabilities. As stated prior (in section 2.1), simply setting the command gains of individual cascaded filters to the intended level will result in a drastically different output due to interaction between adjacent filters [1]. To combat this cascade graphic equalisers have aimed to intake the command gains and have separate internal gain values that will produce the desired output. This idea was first implemented by Abel and Berners [12], which was later improved upon by Oliver and Jot [7]. Both papers fundamentally suggest solving this issue with cascading filters by reducing the problem into a form where a set of linear equations need to solved and optimized for [1]. The later study [7] demonstrates the gain values for which the linear equations need to be solved for (seen below in a simplified example case Eq.1) and how this looks graphically (seen in Fig.3).

$$G_1 = p_1(c_1) + p_2(c_1) + p_3(c1) = g_{11} + g_{21} + g_{31} \tag{1}$$

Essentially the desired gains will therefore be produced by solving the linear set of equations at each specific frequency point so that the summation of each relevant filter gain is equal to it. This becomes more complicated as the amount of overlapping filter sections increases, but by manipulating the matrix algebra a final form can be found for the new correction gain k_n with the addition of a diagonal prototype gain P that remains constant when calculating the response matrix [7]. Their proposed

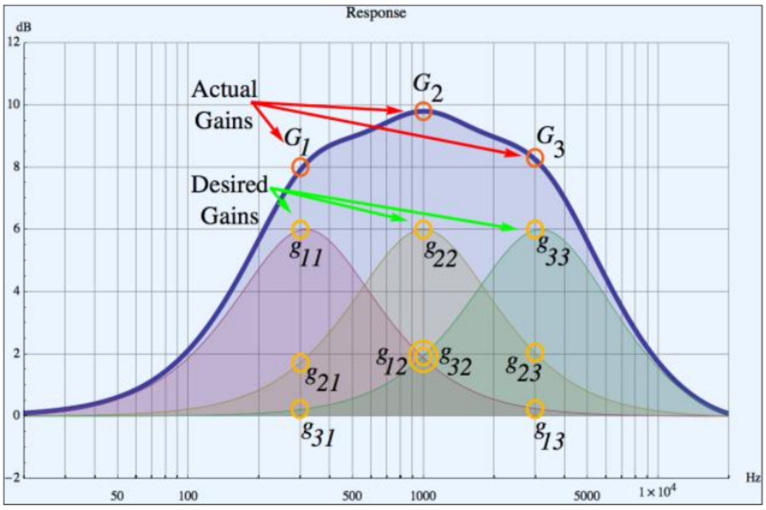

Figure 3: *Cascaded gain summation [7]*

filter also has a separate procedure, median gain offset, for producing a perfectly flat response that is usually very difficult to achieve with a cascade structure, due to the nature of each individual filter having a moderately sharp Q/resonance. Overall, the maximum error produced by this method is 2.28dB in their test cases, when creating a 10-band equaliser [7].

Using this fundamental structure, Välimäki and Liski [13] developed a more accurate 10-band cascaded graphic equaliser by specifying a neighbouring centre frequency gain point, cg_m (note the user can select it, although the paper states that a value of $c = 0.3$ produced the best results), which then defines the bandwidth of each individual filter. Slight variation of these bandwidths is required in the higher frequencies due to them becoming asymmetric, which the authors adjusted separately [13]. This bandwidth definition in the neighbouring filters helps minimize variance with the filter overlap, but needed further adjustments to have a hi-fidelity audio error maximum (defined as $< \pm 1$dB). This involved further refinement of the prototype gain values and an addition of an extra intermediate frequency point between the command gains. This addition involved an additional update of the prototype gains but resulted in a 10-band equaliser that stays under the hi-fi error limit throughout multiple difficult tests (zigzag, every third down and O&J [7] worse case), while also having less operations than previous accurate methods. When adding the additional frequency points the interaction matrix becomes non-square (19-by-10 for the 10-band case), which requires a pseudo-inverse operation that can then be optimized using a least-squares (LS) method with only one iteration since further iterations did not provide significant improvements [14].

A further development of this method was also performed by Välimäki and Liski [14], in which they increased the size of the equaliser so that it became a third-octave equaliser. Therefore, the design now has an interaction matrix which is 61-by-31 (due to centre and geometric mean frequency points), where the highest six bands need to be adjusted by hand to account for the asymmetry in their shape. An important change to note is with the bandwidth definition level is now $c = 0.4$, while the initial prototype gain of $g_p = 17\text{dB}$ remained the same from previous octave equaliser case. This equaliser produced results similar performance to the octave equaliser it originated from, maintaining the $\pm 1\text{dB}$ tolerance defined for high quality audio applications but with a higher resolution in frequency. The entire filter magnitude response can be visualised in Fig.4 below.

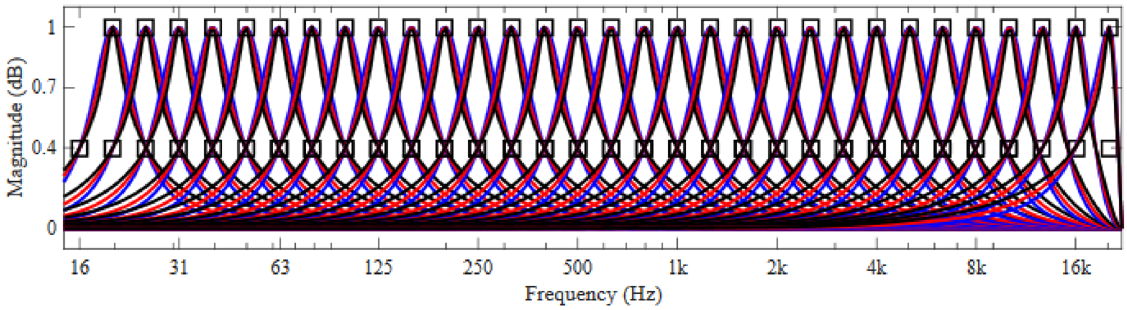

Figure 4: *Filter visualisation for third-octave band cascade graphic equaliser [14]*

Moving onto parallel graphic equalisers, one of the best performing methods is based on the fixed pole IIR filters (described in section 2.2) and was developed further by Rämö et al. [5]. The key element of using this filter design is that not only are the poles able to be selected appropriately for the audio spectrum but also once these poles are selected related to the amount of command gains the only optimization/-computations that will be necessary will be the numerator coefficients and bypass gain d_0. In the investigation, for a third-octave equaliser design it was a sufficient compromise between accuracy and computation speed to have twice as many poles as command gains to accurately form the equalisation. However, this structure struggles when the levels of the equaliser are negative on the logarithmic scale, i.e. small in linear terms, because the LS optimization of the leftover parameters is done on a linear scale. To combat this a weighting function is applied to scale the target and filter response accordingly so that an equal LS error is achieved for any type of target magnitude response. The downside of improving the accuracy with this weighting is that the pseudo inverse matrix used in the optimization task cannot be pre-computed anymore, however depending on the requirements on accuracy and speed the appropriate sub-method of this can be implemented [5].

While this method does increase computation due its complexity, the fact that it can utilize parallel computing with a GPU [9] it partially makes up for this deficit, while also being comprised of many simpler filter sections (seen in Fig.2). Furthermore, the design is non-iterative and with aspects pre-computed the rest of the implementation could well be carried out in real-time audio applications. The accuracy of this architecture also meets the sub 1dB goal as was the case with the cascade design mentioned previously [13][14].

Furthermore, it is relevant to note that in [14], a comparison of the parallel structure to the third-octave cascade was performed which found that while both stay under

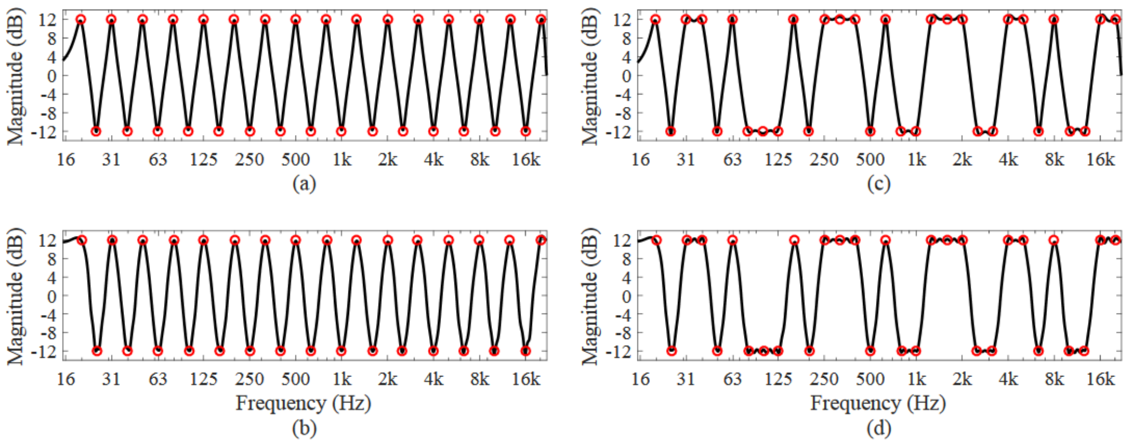

Figure 5: Cascade (a,c) and parallel (b,d) graphic equaliser comparison [14]: A&B ±12dB zigzag and C&D difficult case

the designated error limit of ±1dB their performance varies depending on the target equalisation spectrum. When a ±12dB zigzag was tested the cascade provided a lower maximum error, while the parallel structure proved more accurate in the scenario identified by Oliver and Jot [7] to be difficult for the cascade. Furthermore, when comparing the computation speeds of the two methods, they found that the cascade structure outperforms the parallel in both real-time operations and command gain update times [14]. The accuracy of the two methods can be seen above in Fig.5 from the study.

4 Hybrid Investigation

Having now covered the fundamental structures and high performance techniques for both cascaded and parallel graphic equalisers, their overall similarities and differences can be compared. Both filter architectures suffer from the inherent flaw that adjacent filters will interact with each other. Modern research has shown that for both cascade [13][14] and parallel [5] a high degree of accuracy can be achieved, in both cases a sub ±1dB error was successfully achieved.

When considering design complexity the cascade structure becomes simpler due to two key features: not needing to the consider overall phase of the system, because it will inherit the individual filter phase characteristics, and high performance can be achieved with a smaller architecture [14]. While the accurate cascade topology does require a single optimization iteration to improve the accuracy of the system (after the inclusion of the intermediate frequency points) it still is smaller than its parallel counter-part [14]. Furthermore, both architectures use a similar complexity of individual filters, Orfanidis bi-quad IIR filters for cascade and simplified second order IIR filters with fixed poles for parallel. However, since the cascade has less filter stages and a simpler parameter optimization procedure it becomes faster with respect to both operations and update times [14]. One thing to consider for the benefit of the parallel structure that was stated earlier is that due to its structure it can utilize GPU parallel computing [9], which does improve the efficiency once implemented. In addition to this, while the high performance comparison [14] showed that both could produce high accuracy, an inherent benefit of parallel is that the output will not accumulate quantization noise due to there only being one stage of noise build up [1].

This begs the question whether a benefit/benefits could be gained by formulating the graphic equaliser with a combination of cascaded parallel sections, i.e. a hybrid structure. For example, the larger parallel structure could be segmented into smaller bandwidth filter stages that are then cascaded together. This idea is not entirely new, and there have been implementations with the use of more traditional electronics. One of these was investigated by Adams in 1980 [15] where he aimed to design an automated equalizer to correct for insufficient responses from audio devices. The design utilized a hybrid structure to implement its 10-band equaliser in two stages, for which the key design goal that lead to this was to minimize band interaction but also reduce the amount of noise from the filter sections [15]. A later approach by Erne and Heidelberger [16] is more akin to modern DSP in its approach having the entire architecture implemented digitally. Each filter was implemented as a second order bandpass IIR filter but the key element of the design was the segmentation of the structure into parallel and series to form a 27-band equaliser with minimum phase characteristics. In it, the four cascaded parallel filter banks had centre frequencies selected so that each parallel bank had a larger spacing between adjacent bands, this distribution can be seen below in table 2. This frequency spacing aimed to reduce the band interaction problems of the parallel structure while the division into four cascaded sections would make the computation lighter. Furthermore, the authors noted that the structure could be run on separate devices chained together to fasten processing [16].

Table 2: 27-band hybrid equaliser centre frequency and filterbank division [16]

Filterbank Number	f_{c1} (Hz)	f_{c2} (Hz)	f_{c3} (Hz)	f_{c4} (Hz)	f_{c5} (Hz)	f_{c6} (Hz)	f_{c7} (Hz)
#1	40	100	250	630	1.6k	4k	10k
#2	50	125	315	800	2k	5k	12.5k
#3	63	160	400	1k	2.5k	6.3k	16k
#4	80	200	500	1.25k	3.15k	8k	-

To investigate the feasibility of the hybrid structure in the current state of graphic equalisers using DSP, this filter structure would be tested using the parallel architecture of Rämö et al. [5] and then cascading it using Välimäki and Liski [13]. The first stage of testing this would be to implement the division of the whole filter architecture, into cascaded stages with a similar method as in Erne and Heidelberger [16], for which the division was done as seen in Table 2. The MATLAB code used in [5] is available online [10], which is the basis for this pursuit. The first step of this hybrid adaption is to adjust the target equaliser frequencies into the four separate filter banks, which was performed simply by having each single filter take every fourth centre frequency from the original parallel array. From this each filter would also take be assigned the appropriate gains and pole frequencies to match the equalisation sub-division. In order to avoid any of the newly defined filters not having a target at both 10Hz and at Nyquist, these are added separately. However, these fixes still do not entirely solve the problem caused by these modifications, as can be seen in the Fig.6 below.

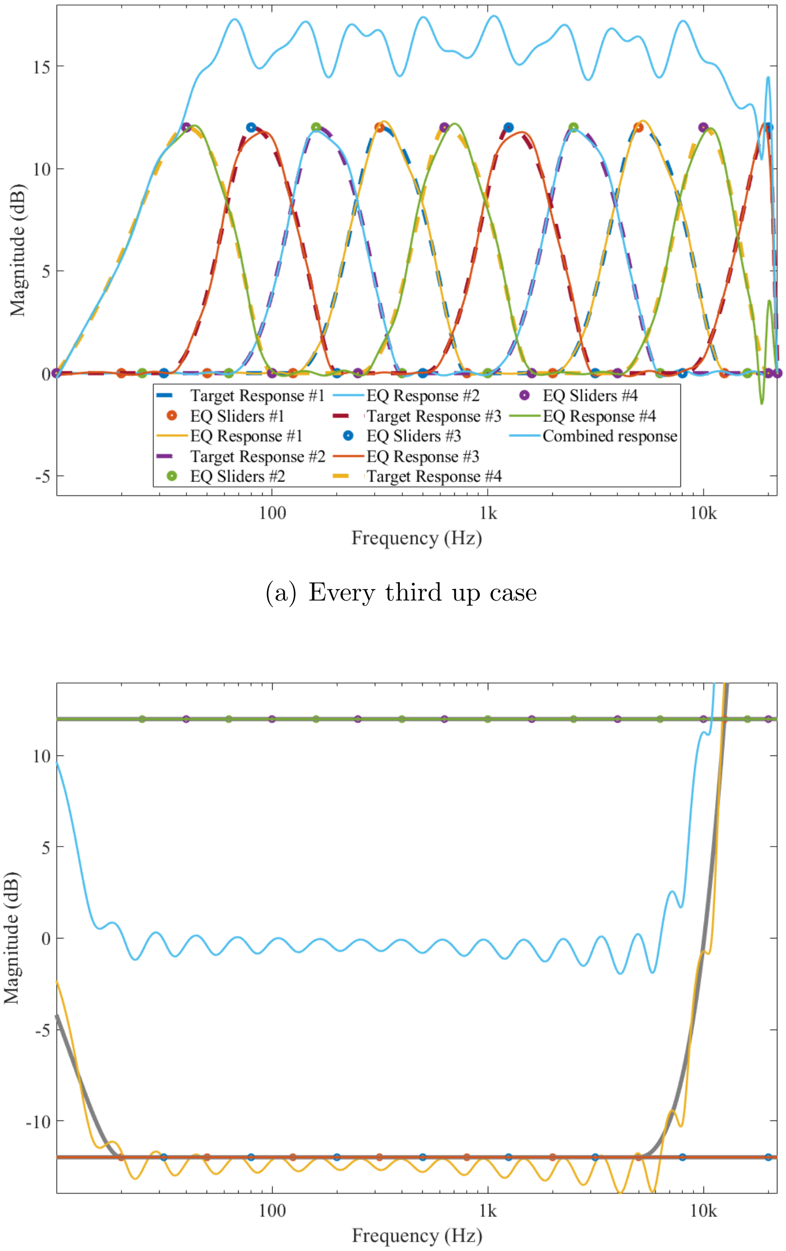

(a) Every third up case

(b) Zigzag case

Figure 6: *Hybrid equaliser initial testing using [5][10] parallel method*

As evident in the every third up case Fig.6(a) the ideally combined response (by summing the outputs of each stage in dB) is far from what is produced when just using the parallel method in [5]. The main issue arises from the fact that the parallel architecture can convincingly follow demanding command points, but now that the filters have been given a less dense array of target frequencies their output also becomes less precise. This leads to an output similar to what was produced by a purely analog graphic equaliser measured in [5]. An important aspect of taking the approach in Erne and Heidelberger [16] with these current high accuracy methods is that solving the issue of band interaction by spacing each filter out is not an issue with either the cascade [13] or parallel [5] method. Furthermore, trying to do so with the parallel method in this case ruins the possible accuracy of the architecture since it does not operate on simple fixed Q filters but rather optimizes to produce the target command gains as accurately as possible. In this case the optimization was performed for each cascade section, but not for the entire spectrum leading to the large errors. While parts of this accuracy problem could be solved by adjusting the resolution of the filters, this would counteract the optimization that could be achieved by introducing the hybrid format to begin with.

These problems are only compounded in the zigzag target case Fig.6(b), in which the resolution of the individual filters results in the zigzag not being formed at all. Rather, the resulting output (in teal as in the previous case 6(a)) has three of the filters holding a constant ±12dB and one filter that changes its magnitude twice over the entire spectrum (since it starts at 0dB and goes to +12dB). This shows how the problems addressed in [16] cannot directly be applied to these modern methods, but rather the hybrid approach should be implemented in another way, by adjusting how the division is performed, how the modern methods are applied or to solve another issue entirely.

One relevant issue that concerns both the parallel and cascade method is computational requirements, which was a key part in their evaluation in section 3 and also in [14]. For the parallel case there has been development by Bank et al. [17] in which the parallel method's accuracy was not compromised but a drastic improvement in computation was achieved. Continuation of these optimization investigations with a hybrid structure as the foundation could utilize it to obtain further computational cost refinement.

5 Conclusions and Future Outlook

This paper has covered the two fundamental equaliser structures, parallel and cascade, which form the graphic equaliser using DSP. While both methods have their benefits, each also has their downsides. Simplicity of design and speed benefit the cascade, while simultaneous computation and accuracy benefit the parallel method. Earlier studies have proposed a possible combination of the two methods to reduce band interaction and provide improvements in performance. An approach shown by Erne and Heidelberger [16] was adapted to modern DSP methods in graphic equalisation to see if it could provide benefits, however the adaption fundamentally aimed to solve a problem that was not present anymore in either cascade or parallel architectures and simply worsened accuracy. Adapting the parallel method directly into a hybrid structure did not correctly implement the possible improvements of the system, because it aimed to accurately produce specified command gains through optimization but the cascaded sections are optimized separately in this investigation leading to a poor end response (see Fig.6). Therefore, further testing of the hybrid system would be needed for example by adjusting the current parallel method [5] to perform optimally within the new architecture. Testing other less accurate parallel methods could also be investigated to see whether the hybridization provides improvements, while also reducing computational cost by simplifying the underlying method.

Modern studies [5][13][14] have provided highly developed versions of both parallel and cascade structures that produce extremely accurate graphic equalisers. These designs are complex to achieve their accuracy (sub $\pm$1dB error), but are both still aiming to reduce computational cost without losing this achieved precision [17]. Therefore, the hybrid structure could also be investigated to reduce computational costs of these modern methods. For example, if the hybrid layout could be able to provide parallel computing possibilities by adjusting the structure of the cascade structure shown by Välimäki and Liski [13][14]. Conversely the problem could be approached in the opposite manner and provide a reduction in computation for the parallel method [5] by formulating it in cascaded sections. Since the accuracy of both cascade and parallel methods have drastically improved with modern DSP, this direction could be an effective use of a hybrid equaliser structure.

References

[1] J.D. Reiss & V. Välimäki. All about equalization: solutions and frontiers. In *Applied Sciences*, volume 6, page 129, May 2016. Accessed on 15.04.2019: https://www.mdpi.com/2076-3417/6/5/129.

[2] D.T.N. Williamson. Design of tone controls and auxiliary gramophone circuits. *Wireless World*, (55):20–49, 1949.

[3] P.J. Baxandall. Negative-feedback tone control. *Wireless World*, (58):402–405, 1952.

[4] G. Massenburg. Parametric equalization. In *Convention of the Audio Engineering Society*, Los Angeles, CA, USA, 2-5 May 1972.

[5] J. Rämo; V. Välimäki; B. Bank. High-precision parallel graphic equalizer. *IEEE/ACM Transactions on Audio, Speech and Language Processing*, 22(12):1894–1904, December 2014.

[6] ISO. ISO 266, Acoustics—preferred frequencies for measurements, 1975.

[7] R.J. Oliver & J.M. Jot. Efficient multi-band digital audio graphic equalizer with accurate frequency response control. In *Convention of the Audio Engineering Society*, New York, NY, USA, 29 September - 1 October 2015. Accessed on 15.04.2019: http://www.aes.org/e-lib/inst/browse.cfm?elib=17963.

[8] B. Bank. Direct design of parallel second-order filters for instrument body modelling. In *Proceedings of International Computer Music Conference*, pages 458–465, Copenhagen, Denmark, 4-9 May 2007. URL: http://www.acoustics.hut.fi/go/icmc07-parfilt.

[9] J.A. Belloch; B. Bank; L. Savioja; A. Gonzalez; V. Välimäki. Multi-channel IIR filtering of audio signals using a gpu. In *IEEE International Conference Acoustics, Speech and Signal Processing (ICASSP)*, pages 6692–6696, Florence, Italy, 4-9 May 2014. Accessed on 15.04.2019: https://ieeexplore.ieee.org/document/6854895.

[10] B. Bank. Transfer function modeling and equalization by fixed-pole parallel filters, 23 May 2018. Link: http://home.mit.bme.hu/ bank/parfilt/.

[11] M. Virgulti; S. Cecchi; E. Piazza. IIR filter approximation of an innovative digital audio equalizer. In *Proceedings of International Symposium on Image and Signal Processing and Analysis (ISPA)*, pages 410–415, Trieste, Italy, 4-6 September 2013.

[12] J.S. Abel & D.P. Berners. Filter design using second-order peaking and shelving sections. In *Proceedings of the International Computer Music Conference*, Coral Gables, FL, USA, 1-6 November 2004.

[13] V. Välimäki & J. Liski. Accuracte cascade graphic equalizer. *IEEE Signal Processing Letters*, 24(2):176–180, 2017. https://doi.org/10.1109/LSP.2016.2645280.

[14] V. Välimäki & J. Liski. The quest for the best graphic equalizer. In *Proceedings of the 20th International Conference on Digital Audio Effects (DAFx-17)*, pages 95–102, Edinburgh, UK, 5-9 September 2017. Permalink: http://urn.fi/URN:NBN:fi:aalto-201812106225.

[15] R.W. Adams. An automatic equalizer/analyzer. In *Proceedings of 67th Audio Engineering Society Convention*, New York, NY, USA, 1 October - 3 November 1980. Permalink: http://www.aes.org/e-lib/browse.cfm?elib=3696.

[16] M. Erne & C. Heidelberger. Design of a DSP-based 27 band digital equalizer. In *Proceedings of 90th Audio Engineering Society Convention*, Paris, France, 19-22 February 1991. Permalink: http://www.aes.org/e-lib/browse.cfm?elib=5664.

[17] B. Bank; J.A. Belloch; V. Välimäki. Efficient design of parallel graphic equalizer. *Journal of the Audio Engineering Society*, 65(11):817–825, 2017. Permalink: http://www.aes.org/e-lib/browse.cfm?elib=19355.