Audio Engineering Society
Convention Paper 6577
Presented at the 119th Convention
2005 October 7–10 New York, New York USA

This convention paper has been reproduced from the author's advance manuscript, without editing, corrections, or consideration by the Review Board. The AES takes no responsibility for the contents. Additional papers may be obtained by sending request and remittance to Audio Engineering Society, 60 East 42nd Street, New York, New York 10165-2520, USA; also see www.aes.org. All rights reserved. Reproduction of this paper, or any portion thereof, is not permitted without direct permission from the Journal of the Audio Engineering Society.

An ultra high performance DAC with controlled time domain response

Paul Lesso[1] and Anthony Magrath[2]

Wolfson Microelectronics plc, Westfield House, 26 Westfield Road, Edinburgh, EH11 2QB
[1] Paul.Lesso@wolfsonmicro.com
[2] Anthony.Magrath@wolfsonmicro.com

ABSTRACT

This paper describes the design of an ultra-high performance stereo digital-to-analogue converter (DAC) employing advanced digital filtering techniques. Recently there has been a renewed interest in the time-domain properties of digital filters used for interpolation and decimation. Linear phase FIR filters, which have proliferated digital filter design for the last two decades, have the undesirable properties of pre-ringing and high group delay. Conversely, minimum phase filter filters, which offer lower levels of pre-ringing, do not have a uniform phase response. This paper describes the trade-offs in the design of filters with controlled pre-ringing, coupled with desirable phase and magnitude characteristics. The paper also describes architectural choices in the implementation of the DAC signal processing chain, required to achieve commensurate analogue performance.

1. INTRODUCTION

This paper reports on the architectural tradeoffs made in the design of a high performance noise shaped DAC. The DAC supports multiple filters for the 8x oversampling stage, providing a choice of time-domain and frequency-domain characteristics.

Conventional audio filter designs tend to focus on the frequency-domain response and neglect the time-domain response. However there is now an increasing interest in the effect that time-domain properties of these filters have on the perceived audio quality. Some properties, such as a large pre- and post-echo due to excessive passband ripple are well known, but others such as pre-ringing have received less attention. Traditionally, linear-phase filters have been used to avoid group-delay distortion, however this comes at the cost of pre-ringing in the time domain which is widely believed to less than optimal for listening. Minimum-phase filters are an alternative that reduces the pre-ringing at the cost of some group delay distortion and increased post-ringing. Recent work by Craven also highlights the possibility of using apodising filters to reduce distortions introduced elsewhere in the chain. This paper explores some of the trade-offs between minimum-phase filters and linear-phase filters for a high performance DAC. We also look at the advantages/disadvantages of filters with a prescribed passband droop.

2. INTERPOLATION FILTER DESIGN

In this section we discuss the relevant metrics for the design of a series of interpolation filters.

2.1. Passband ripple

It is well known that excessive passband ripple leads to pre and post echoes [4]. Maxflat filters [16] offer much less ripple, yet 'waste' zeros ensuring no ripple over that passband. In the designs presented here, pre- and post-echo are at least 100dB down from the main impulse.

2.2. Stop-band attenuation

More stopband attenuation is desirable for image suppression in the filter chain, but the deeper the stopband the more ringing in the filter. For example, a filter design using the Kaiser windowing method [2] has a length N given by

$$N = \frac{attn - 8}{14.36\Delta}$$

where *attn* is the stopband attenuation in dB and Δ is a function of the transition band. Whilst this expression does not represent an optimal filter design (and indeed Remez-designed filters follow a more complex rule) , it does show that (to first order) the number of taps in the filter and hence the pre-ringing of symmetric filters is directly proportional to the stopband attenuation. This indicates that the more aggressive the stopband the more ringing in the filter and the larger the delay through the filter. This suggests that increasing the stopband may be a law of diminishing returns. Conversely the minimum phase filters described below may offer a way to overcome this.

2.3. Anti-alias and apodising filters

Black suggests that transition band aliasing is a source of distortion [5] in playback systems. Certainly there is anecdotal evidence to suggest that devices that are fully attenuating by Fs/2 sound better than those that allow some aliasing. For a 44.1kHz sampling rate large filters are required in order to fully attenuate by Fs/2 which as discussed above, can lead to unacceptable levels of pre-ringing.

Craven recently introduced the concept of an apodising filter: a filter that can be used to control the time smear of a whole recording and reproducing chain [12]. This type of filter can reduce the pre- and post-ringing of the impulse response; and indeed it has been shown that if correctly designed they can compensate for any number of filters.

The key to designing an apodising filter is to ensure that the response of the apodising filter is zero in the transition bands of the other filters that it is compensating for.

Since an apodising filter requires the cutoff point to happen before the Nyquist frequency these filters can also be used to eliminate transition band aliasing, even if it has occurred in the sampling ADC[13].

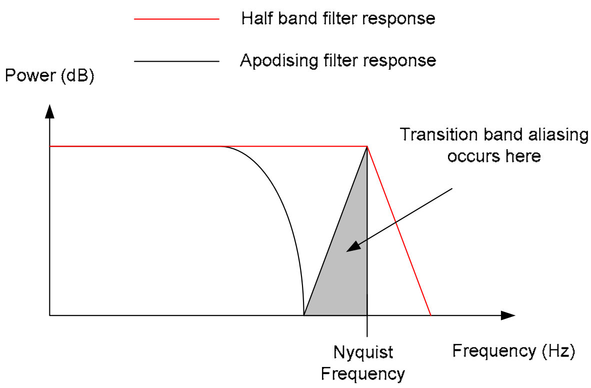

Figure 1 - Frequency response of an apodising filter

As shown in Figure 1, if half-band filters are used in the ADC then transition band aliasing can occur (denoted by the gray shaded area). By placing the cutoff frequency of the filter below this frequency we can remove any noise in the region and compensate for both the time dispersion of the system and transition band aliasing from the ADC. Note that apodising filters that eliminate transition-band aliasing are really only possible at higher sample rates.

2.4. Linear phase versus minimum phase filters

It was first suggested by Lagadec [1] that pre-ringing might be a source of audible distortion from converters. It is also consistent with our knowledge of hearing that a pre-response is more significant than the post-response (*q.v.* the Haas effect and temporal masking [10]). Temporal masking occurs when a sudden stimulus obscures sounds preceding (backwards masking) or following (forwards masking) the stimulus. The human ear has much more forward masking than backward masking and this suggests we are less sensitive to post-echo than we are to pre-echo.

It has been shown that the human ear has a inherently non-linear characteristic [6] and Lesurf has shown that effect of pre-ringing on a non-linear system representative of the human ear is to produce an effect which may be audible [8]. Pre-ringing rarely occurs in nature – as pre-ringing corresponds to hearing the effect of the sound source before the originating sound reaches the listener. However in some natural cases this does occur – for example when sound is picked up from transmission though the ground before the it is heard from the loudspeaker [9].

Hence it is interesting to consider the degree to which we can reduce pre-ringing. Linear phase filters are widely used because they introduce no group-delay distortion but at the cost of pre-ringing. Conversely minimum phase filters have some group-delay distortion but little pre-ringing. Figure 2 shows the difference in the impulse responses for two such filters, designed for 120dB stopband attenuation.

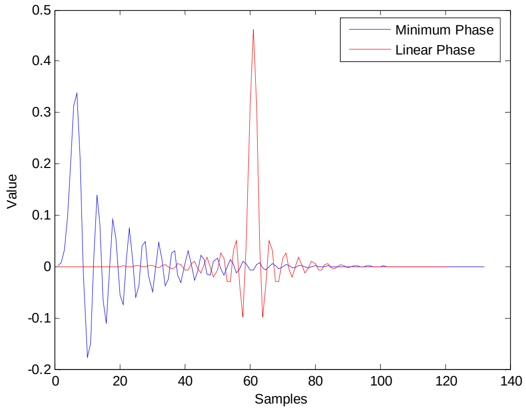

Figure 2 -Impulse response of linear and minimum phase filters

One filter is linear phase and hence has the main lobe in the middle of the impulse; conversely the minimum phase filter has the main lobe much earlier and therefore has a much lower latency.

2.5. Minimum phase filters

At this point it is worth defining minimum-phase and looking at how these filters can be designed.

A minimum phase filter is a filter that has all the poles and zeros inside the unit circle. This leads to the interesting property for testing if a filter transfer function *H(z)* is minimum phase :

> *H(z)* is minimum phase if *H(z)* and *1/H(z)* are both stable

2.5.1. Design of Minimum Phase Filters

It is also worth noting that any non-minimum phase filter can be factorized into an all-pass filter and a minimum phase filter. This can be easily shown – consider a non-minimum phase filter T(z) which has a single zero, *a*, outside the unit circle

$$T(z) = H(z)(1 - az^{-1})$$

hence

$$T(z) = H(z)\left(1 - \frac{z^{-1}}{a}\right)\left(\frac{1 - az^{-1}}{1 - \frac{z^{-1}}{a}}\right)$$

where

$$H_{MINIMUM-PHASE} = H(z)\left(1 - \frac{z^{-1}}{a}\right)$$

and

$$H_{ALL-PASS}(z) = \left(\frac{1 - az^{-1}}{1 - \frac{z^{-1}}{a}} \right)$$

Obviously the above argument holds for any number of non-minimum phase poles and zeros. This represents one possible method to design a minimum phase filter – design a linear phase equivalent and then factor out all the non-minimum phase terms. This method is known as zero inversion, and results in a minimum-phase characteristic without affecting the magnitude response. However this method can lead to the creation of double zeros and hence the creation of a sub-optimal filter. This method also tends to be unsuitable for the synthesis of long filters with deep stopband attenuation due to numerical problems.

Other methods include spectral factorization which involves generating a filter with the transfer function of the form $\left|H(z)\right|^2$ where H(z) is the desired response. The zeros outside the unit circle and double zeros on the unit circle are deleted and then the polynomial is converted back to coefficients. However this method is prone to numerical errors when converting from coefficients to polynomial form for long filter lengths. Also some filter algorithms, such as Remez, might struggle to design a filter with the response $\left|H(z)\right|^2$ if the stop band is deep. E.g. for filters with a 120dB stop band we would need to design a prototype filter with a 240dB stop band.

There is not the scope for a full discussion of design methods for minimum phase filters here but the authors note that that for filters with deep stopband attenuation there are a host of numerical issues to overcome [3]. A variety of methods exist to perform this spectral factorization. However it must be noted that these methods do not give direct control over the impulse response.

For low-order filters genetic algorithm based designs can achieve very successful results with a large degree of control over the time and frequency response but we found such methods too slow to use for designing filters with -120dB stopband.

2.6. Audibility of group delay distortion

Research has shown the ear is relatively insensitive to group delay distortion of several milliseconds for low frequencies (<1kHz) and insensitive to +/-0.5ms over the 1-5kHz band [19]. Other work shows that the sensitivity to group delay distortion falls after 4kHz and therefore group delay distortion in the upper regions of the audio band is much less audible [20].

For a typical minimum phase filter designed for 44.1kHz the group delay distortion up to 10kHz is under 2 samples (less then 46µs) and may be inaudible.

2.7. Interpolation filters for 44.1kHz and 48kHz audio

Since there is no definitive indication that minimum phase filters sound better than linear phase filters, particularly given the phase distortion over the audio band it was decided to implement a series of filters and allow the end user to decide which filters to use. The reduced latency of the minimum phase filters may be ideal for studio situations where a low latency is needed. Conversely some users may prefer the standard linear phase filters.

The filters we chose to implement are:

1. linear phase
2. minimum phase with tailored response
3. minimum phase
4. no-alias linear phase
5. no-alias minimum phase

2.7.1. Linear phase filter

This is the standard filter built using half-band architecture for backward compatibility. These filters would be used for a brickwall filter when there is an apodising filter elsewhere in the chain.

2.7.2. Minimum phase with tailored response

At 44.1kHz this filter has a rippled passband to 18.375kHz and then a smooth droop of 0.4dB to 20kHz. This filter has better attenuation at Fs/2 than a half band filter (-28dB compared to -6dB) but more droop over the passband. It is intended to be a compromise between filter 1 and filters 4 /5 above.

As can be seen in Figure 3, the difference in group delay up to 10kHz is less than 1 sample and it rises sharply after this up to a difference of 14 samples at 20kHz.

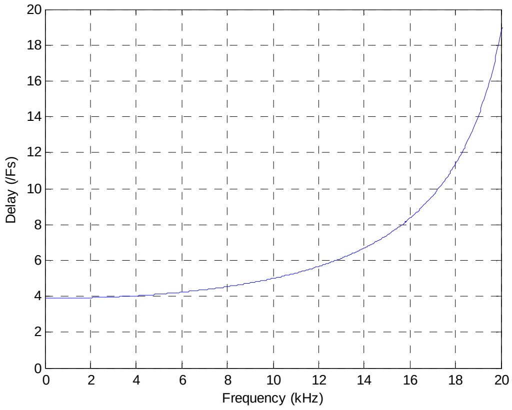

Figure 3 - Group delay of filter 2 for 44.1kHz

2.7.3. Minimum phase filter

This filter was designed using the methods described in section 2.5.1 and has a low latency at the cost of some group delay distortion. The group delay distortion is similar to that of the filter above but the passband is flat to 20kHz and hence there is less attenuation at Fs/2.

2.7.4. No-alias linear and minimum phase filter

These filters were designed to be fully attenuating at Fs/2. They have a considerable droop at 20kHz when Fs=44.1kHz. They have the same frequency domain response but as can be seen from figure 4 they have very different time domain responses. Note all the impulse responses beyond this point in the paper show the power of the impulse response on a logarithmic scale.

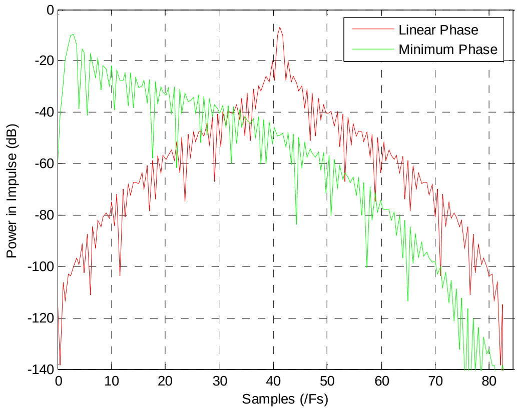

Figure 4 - Impulse response of filters a minimum and linear phase filter

3. HIGH SAMPLE RATE AUDIO

The advantages of high sample rate audio are widely debated. Several people have shown that the advantage of the higher sample-rate audio is not the extended bandwidth above 20kHz (even though there is some evidence that power above 20kHz has some effect [7]) but is instead the fact that we can use the extended bandwidth to tailor the transition band and reduce the time dispersion of the impulse response [11][13].

It has been argued that the advantage of DSD is that the time domain dispersion of the impulse response of the system is limited and so equivalently it makes sense to try to limit the extent of the impulse response of the filters discussed here.

Even at the higher sample rates there are trade-offs to be made. Figure 5 shows the impulse response of two filters for a 96kHz input, one with a cutoff at 20kHz and the other with a cutoff at 40kHz. The impulse response for the filter with the cutoff at 20kHz is under half the width and has considerably less pre-ringing[1]. The extra degree of freedom at higher sample rates therefore makes it possible to design much more interesting filters. Equivalently, the 40kHz filter has the same impulse response of a filter running at 48kHz, with 20kHz cutoff. This shows that at higher sample-rates it is possible to design filters that have substantially lower time dispersion, and perhaps this explains the reported

[1] Note the impulses are time aligned to show the difference in the time dispersion and latency.

improvement in audio quality associated with higher sample-rate audio [15].

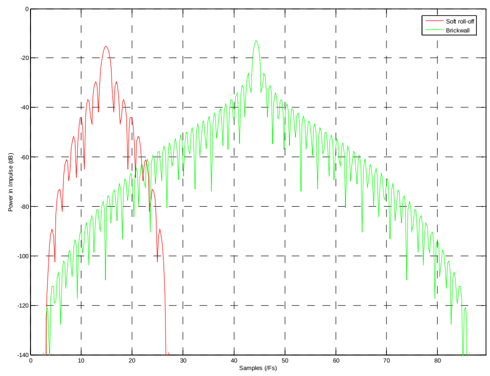

Figure 5 - Impulse response for 20kHz cutoff and 40kHz cutoff for 96kHz interpolation filter.

3.1.1. Interpolation filters for high sample rate audio

For the higher sample rates we settled on 5 filters

1. linear phase 'soft knee filter'
2. minimum phase 'soft knee filter'
3. linear phase brickwall filter
4. minimum phase apodising filter
5. linear phase apodising filter

3.1.2. 'Soft knee' filters

These filters take advantage of the larger transition band to reduce the dispersion and delay through the filter, such as shown in Figure 5. These filters are fully attenuating by Fs/2 but have a passband that only extends to 20kHz. Figure 6 shows the group delay through the minimum-phase soft-knee filter, which has a variation of around 1 sample up to 15kHz, and less than 2 samples over the full audio band.

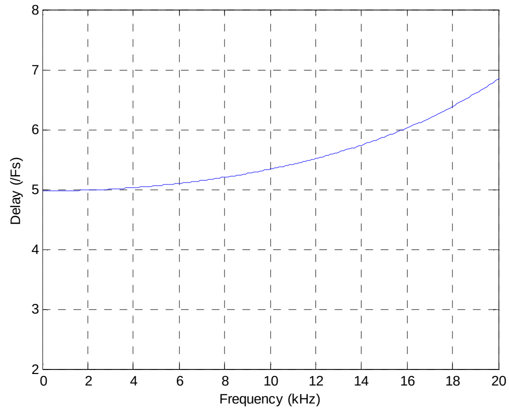

Figure 6 - Group delay variation for a minimum phase filter at 96kHz

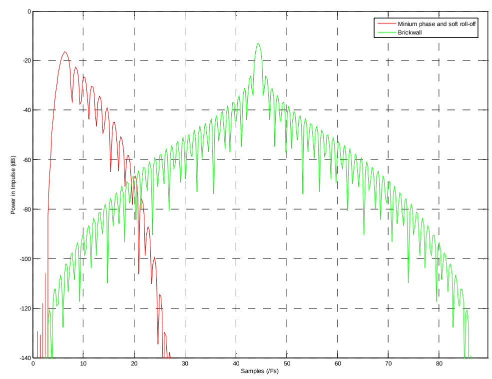

Figure 7 - Comparison of impulse responses for minimum and linear phase apodising filters

3.1.3. Brick wall filters

These filters have a passband that scales with frequency thus having a small transition band and a large amount of ringing. It is intended that these filters should be used in a system which already has an apodising filter in the chain.

3.1.4. Apodising filters

As discussed above, these filters sacrifice some the transition band to attenuate out of band products and act to reduce the ringing of the impulse response of the whole system. Figure 7 shows the impulse response for the minimum-phase and linear-phase apodising filters. Note the low group delay and minimum pre-ringing in the minimum-phase filter.

4. INTERPOLATION ABOVE 8FS

Following interpolation to 8Fs, the signal requires up-sampling to the final DAC output rate of 256Fs. This is performed using a 3rd order Cascaded Integrator Comb (CIC) filter. The purpose of this filter is to adequately attenuate all images from 8Fs up to 128Fs. CIC filters offer an extremely compact way of achieving high attenuation at multiples of the input rate - in this design, all images are attenuated by 70dB for an input band-limited to 20kHz, with a hardware cost of only six adders. The reason for providing such a high level of attenuation is to reduce susceptibility to clock jitter [25], since any out-of-band periodic components may cause baseband intermodulation products in the presence of clock jitter.

Figure 8 – CIC Filter

5. FILTER QUANTIZATION

A problem often overlooked until the last stage of filter design is the problem of filter quantization. When the coefficients are initially designed they will be specified to floating point accuracy and then must be converted to fixed point numbers for actual use. This will introduce passband ripple and reduce the stopband. If the targets are not met there are two options – either redesign the filter hardware or try to optimize the coefficients. A simple quantization to the nearest value will not represent an optimal solution.

One method is to use noise shaping when the coefficients are quantized [22]. Work by the authors has shown that this can achieve impressive results, but the full advantage of noise shaping cannot be taken into account since noise shaping only works over infinite sequences. There are proprietary methods such as IVQ[23] that has been design to overcome this limit.

Genetic algorithms (GAs) are also widely used due to the high degree of non-linearity between coefficient quantization and the target cost function. However most GAs have a series of options, which make a big difference to convergence leading to the need for a meta-optimizer to optimize the control of the optimizer. This method is widely used because it can produce very impressive results but it can be very time consuming to run.

For the design of the filters detailed above the authors have used a variant of a bit flipping algorithm previously suggested in [21]. This algorithm is a brute force approach akin to integer-programming methods. In this method each bit of the coefficient in turn is ‘flipped’ (i.e. a 0 becomes 1 and vice-versa) and it is seen if this improves the filter characteristics. Then two bits are flipped simultaneously and so on. In this way we exhaustively search through the solution space to find a more optimal solution. It is possible effectively to reduce the coefficient wordlength by several bits in this manner for a small simulation time. We used this method in preference to the GA method for the speed at which it ran.

6. QUANTIZATION EFFECTS IN DIGITAL SIGNAL PATH

Care must be taken with digital audio to ensure that quantization effects do not dominate the performance of the chip. For a DAC this is less of a concern since the analogue noise floor will hide the digital noise floor. Nevertheless it is important to ensure that the digital signal path is a linear as possible and eliminate quantization effects to an acceptable level.

6.1. DC sweeps

There are two main effects through which quantization error manifests itself – distortion and noise floor modulation. It is possible to analyze these effects by sweeping a DC input from 0dBFS to 144dBFs and measure the residual noise + distortion. For a well designed and properly dithered system the residual should be constant across the whole sweep. This is only achieved for DACs by design [26].

6.2. Dither

Dither is widely used to eliminate both tones and noise floor modulation [14]. However dither comes with a cost to the noise floor.

- 1 LSB rectangular probability density function (PDF) (RPDF) dither raises the noise floor by 3dB and eliminates tones.

- 2 LSB Triangular PDF (TPDF) dither raises the noise floor by 4.77dB and eliminates tones and noise floor modulation.

In addition to the above dither types, High-Pass TPDF (HPTPDF) dither is also supported, which spectrally shapes some of the dither noise out of band. Note that the noise modulation introduced affects only the digital noise floor, which is below the analogue noise floor. However it is not known how far below the analogue noise floor the digital artifacts need to be before they become inaudible. Hence we have designed this DAC to have a benign digital noise floor when TPDF dither is used.

Dither was applied at all major truncations in the signal path. Dither was needed at several places – once in the ALU multiplication, once in the ALU accumulator prior to writing back to RAM, in the CIC interpolator and in the sigma-delta modulator. To ensure that there is no correlation between the various truncations multiple independent dither sequences were generated.

Recent work shows that dithering in a ADC/DAC chain does have an audible effect [18] and that RPDF and TPDF do have audibly different effects under certain conditions.

7. NOISE-SHAPING FOR 120DB+ DACS

Noise-shaping is used to reduce the wordlength at the output of the CIC filter to 6-bits. Due to the high oversampling ratio used, it is possible to use only 2^{nd} order noise-shaping yet still maintain the low noise-floor of the DAC. A little-cited benefit of second-order noise-shaping is that the correlation between adjacent samples is minimized, compared to third or higher-order noise-shaping. For the low-noise baseband properties of the noise-shaping to be preserved through the final DAC output stages, it is essential that both timing and amplitude information is accurately preserved. If information in either domain is lost, the underlying mathematics behind the noise-shaping is disrupted, resulting in an increase in baseband noise. This effect is more apparent for high-order noise-shaping characteristics, where the information is spread across hundreds of samples. For this reason, the use of 2^{nd} order noise-shaping improves immunity to timing errors, e.g. clock jitter; and immunity to amplitude errors, e.g. capacitor mismatch. This immunity is further enhanced by the use of a large number of quantizer bits, which again reduces the impact of timing and amplitude errors.

An added benefit of using only 2^{nd} order noise-shaping is that a higher modulation index is possible, compared to high-order modulators, thus maximizing the dynamic range of the converter. As with the preceding interpolation stages, the quantizer is carefully dithered to ensure the output is tone-free.

8. DYNAMIC ELEMENT MATCHING

A unique DEM scheme is used, as reported previously [26] which allows the use of binary-weighted DAC elements. The key benefit of this scheme is to minimize the number of elements (capacitors) and associated switching and routing overheads. The DEM algorithm shapes the mismatch error between elements so that the audio-band contribution of the error is reduced.

Fundamental to the operation of the DEM are pairs of coupled sigma-delta modulators, which select the elements according to the binary values of the noise-shaper output. The capacitors are either selected with the same polarity, or opposite polarity, according to a decision made by the vector-quantizer which couples the two modulators together.

The sigma-delta modulators produce a strong limit cycle at FS/2, meaning that the outputs tend to oscillate between the two states in the shortest possible time. This causes each capacitor in a pair to be utilized with maximum frequency, causing the capacitor errors to appear as high frequency artifacts, with minimal AC components at low frequencies. Additionally, linearity is improved, since no particular DAC elements are used predominantly for particular input values.

The modulators are dithered to minimize noise-modulation without imparting a significant noise penalty.

9. CONCLUSIONS

The paper details the design considerations for a high performance DAC with controlled time domain response. Particular focus was given to the design issues related to interpolation filters. We discussed practical considerations for the design of filters with reduced time dispersion which have been reported to offer improved sound quality. Filter designs were suggested that offer a compromise between low pre-ringing and low group delay distortion. Additionally the design of high-rate apodising filters were discussed, which allow time dispersion effects in any part of the reproduction chain to be compensated for. We discussed in detailed practical implementation of the filters, including coefficient optimization for finite wordlength implementation, signal quantization and dithering. Furthermore, aspects of noise-shaping and dynamic element matching were discussed which allows the DAC to be geared toward high-end performance with real-world effects such as clock jitter and element mismatch.

10. ACKNOWLEDGEMENTS

The authors thank Chris Travis for his many useful insights into this work.

11. REFERENCES

[1] Lagadec, R. "Dispersive Models for A-to-D and D-to-A Conversion systems", AES 75th Convention, March 27-30,1984. Paris, France. Preprint 2097.

[2] D. Sclichtharle," Digital Filters",Springer-Verlag

[3] N. Damera-Venkata and B. L. Evans, "Design of Optimal Minimum Phase Digital FIR filters Using Discrete Hilbert Transforms", IEEE Transactions on Signal processing, March 2000

[4] J. Dunn, "Anti-alias and anti-image filtering: The benefits of 96kHz sampling rate formats for those who cannot hear above 20kHz", AES 104th Convention, May 1998

[5] R. Black, "Anti-alias filters: the invisible distortion mechanism in digital audio?", AES 106th convention, 1999

[6] V.M. Equiluz, "Essential nonlinearities in Hearing", Physical Review Letters, 29 May 2000

[7] Oohasi et al., "Inaudible High-Frequency Sounds affect Brain Activity: Hypersonic effect", Journal of Neurophysiology, Vol 83, June 2000

[8] J. Lesurf, "Ringing in the ears", Hi-Fi News, September 2004. Also available at http://www.st-and.demon.co.uk/AudioMisc/hearing/ringing/page1.html

[9] M. Gerzon, "Why do equalizers sound different?" Studio Sound, 1990

[10] W. Yost, "Fundamentals of Hearing", 2000, Elsevier Science

[11] M. Story," A Suggested Explanation for (some of) The Audible Difference Between High Sample Rate and Conventional Sample Rate Audio material", dCS Website, September 1997

[12] P.G. Craven, "Controlled pre-response antialias filters for use at 96kHz and 192kHz", 114th AES convention, March 2003

[13] P. G. Craven," Antialias Filters and System Transient Response at High Sample Rates", Journal of Audio Engineering Society, Vol. 52, No.3, 2004

[14] J. Vanderkooy and S. Lipshitz, "Digital Dither: Signal processing with resolution far below the least significant bit", AES 7th Conference

[15] M.N. Harris et al., "Effects in High Sample Rate Audio Material", AES 20th Conference

[16] S. Samadi et al.," Universal Maximally Flat Lowpass FIR systems", IEEE Transactions on Signal Processing, Vol. 48, No. 7, July 2000

[17] R. Wannamaker, "Efficient Generation of Multi-channel Dither Signals", Journal of the Audio Engineering Society, Vol. 52, No.6, 2004 June

[18] P. Kvist," A Listening Test of Dither in Audio Systems", AES 118th Convention, 2005

[19] D. Preis, "Phase distortion and Phase equalization in Audio Signal Processing – A tutorial review", Journal of Audio Engineering Society, Vol 30, No. 11, 1982 November

[20] D. Preis and P.J. Bloom," Perception of phase distortion in anti-alias filters", 75th AES convention, 1983

[21] A. Krukowski and I. Kale, "Two Approaches for Fixed-Point Filter Design: "Bit-Flipping" Algorithm and Constrained Downhill Simplex Method", 5th International Symposium on Signal Processing and its Applications (ISSPA99), Brisbane, Australia, August 22-25th, 1999

[22] J.J. Nielson, " Design of Linear-Phase Direct-From FIR digital Filters with quantized coefficients using error spectrum shaping techniques", IEEE Transactions on Acoustics, speech and signal processing vol. 37, no.7, July 1989

[23] "IVQ method", Alogol Applications Ltd., 53 High Street, Steyning, West Sussex,BN44 3RE, U.K.

[24] Y. Lim *et al.*, "Finite word length FIR filter design using integer programming over a discrete coefficient space", IEEE Transactions on Acoustics, Speech, and Signal Processing, Volume 30, Issue 4, Aug 1982 Page(s):661 – 664

[25] C. Travis and P. Lesso, "Specifying the Jitter Performance of Audio Components", AES 117th Convention, October 2004.

[26] D. Hossack, P. Frith, J. Hayes, A. Jackson, "Design and Evaluation of an Audio DAC with Non-uniformly Weighted Dynamic Element Matching", 16th UK AES Conference, March 2001