24th INTERNATIONAL CONGRESS ON SOUND AND VIBRATION
23–27 July 2017, London

ICSV24
23-27JULY 2017
LONDON CALLING

FULLY DIGITAL IMPLEMENTATION OF A HYBRID FEEDBACK STRUCTURE FOR BROADBAND ACTIVE NOISE CONTROL IN HEADPHONES

Piero Rivera Benois, Patrick Nowak and Udo Zölzer
Helmut-Schmidt-University, Department of Signal Processing and Communications, Hamburg, Germany
email: piero.rivera.benois@hsu-hh.de

Connecting an Internal Model Controller (IMC) and a Minimum Variance Controller (MVC) into a hybrid structure aims to combine the attenuation capabilities of the individual structures without the need of additional microphones or speakers. Moreover, if connected in a particular way, both controllers can be designed independently from each other. Purely digital implementations of such a structure based on standard Sigma-Delta ADC and DAC, and general purpose DSP lead to significant delays in the processing chain that substantially decrease the noise attenuation frequency bandwidth of the MVC. Because of this, a mixture of a digital IMC and an analogue MVC implementation is suggested in the literature. In this paper, a fully digital implementation of the hybrid structure based on successive-approximation-register ADC and DAC, and a low-latency FPGA filtering technique is presented and evaluated with measurements on a real system. The obtained results show that the extremely low analog-to-analog latency of 2.03 μs enables the MVC to reach the performance of an analogue implementation, with a noise attenuation bandwidth of 870 Hz and attenuation levels between 10 dB and 20 dB. This digital implementation offers the opportunity to adjust the filter coefficients in real-time. In addition, the interaction with the IMC introduces the attenuation of periodic signals and a stabilization of the overall system.

Keywords: Broadband feedback ANC, active noise control headphones, low-delay processing, successive-approximation-register, hybrid feedback structure.

1. Introduction

Active Noise Control (ANC) headphones provide the user with an attenuation of the acoustical noise present in his environment. This attenuation is a mixed effect of the passive characteristics of the headphone's construction materials and the active noise control applied to the noise that effectively enters the ear-cups. The passive attenuation is in the practice very effective in the medium-high frequency range. The rest is actively addressed with ANC through the generation of compensating sound pressure levels played together with the user's music through the headphone drive, such that the environmental noise is canceled out by superposition. To achieve this, ANC systems are generally equipped with a reference microphone to measure the incoming noise, and with an error microphone, used to measure the noise that effectively enters the ear-cup and remains after the cancellation.

A typical approach found in ANC headphones is a feedback solution called the Minimum Variance Controller (MVC). Because of its simplicity and sensitivity to delays in the control chain [1], this structure is normally implemented in analogue circuitry [2]. To extend the attenuation capabilities of the system to periodic signals in a broader frequency bandwidth, the MVC is combined with a digital adaptive Internal Model Controller (IMC) with FxLMS [3] to yield a hybrid structure [4][5][6]. It has

been shown in [7], that the controllers can be designed and optimized independently from each other, if they are connected in a certain way.

In this paper a fully digital implementation of the hybrid structure of [7] is presented. Its implementation is based on decoupled successive-approximation-register ADC and DAC units, which have 1 μs conversion delay each, and a low-latency FPGA filtering technique with a latency of 30 ns. Thus, the overall processing delay goes down to 2.03 μs, independent of the sampling rate. Due to the improvement, the digital MVC generates better attenuation and a broader bandwidth than the latest results [8] at the same sampling frequency.

In the following section, the hybrid feedback structure is generally described and its mixed and fully digital implementations are discussed, showing the many advantages of the latter. This is followed by the control and calculation strategies used to reduce the processing latency. Afterwards, the implementation of the MVC on an FPGA platform is described. At the end, the measurement setup is described, the results are evaluated, and conclusions are drawn.

2. Hybrid Feedback Structure

The hybrid feedback structure proposed in [7], which is the base for this work, is presented in a simplified diagram in Fig. 1. In this diagram $W_a(z)$ is the IMC controller, $W_c(z)$ is the MVC

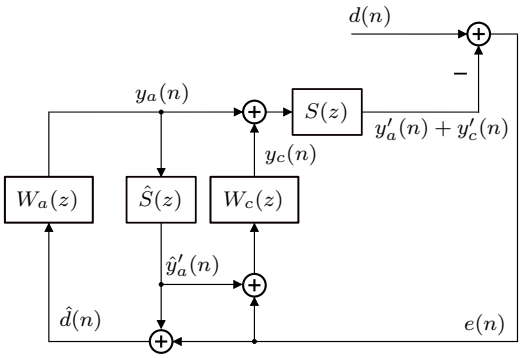

Figure 1: Simplified system diagram of Schumacher's hybrid system

controller, and $S(z)$ comprehends the influence of the headphone's speaker on its input signal and the acoustic path until the place where the error microphone is located. $\hat{S}(z)$ is a model of $S(z)$, that the IMC uses to approximate its control signal at the error microphone's position, $y'_a(n)$. The signal processing starts with $d(n)$, the noise that effectively enters the ear-cup of the headphone, which superposes with the control signals $y'_a(n)$ and $y'_c(n)$ to produce $e(n)$, the error signal. This residual noise is then added with $\hat{y}'_a(n)$ to yield an approximation of $y'_c(n) + d(n)$, the error left only by the MVC. This signal is then filter by both controllers to generate the control signals $y_a(n)$ and $y_c(n)$.

The aspect that differentiates this structure from the one in [4][5][6], is the way the inputs to the MVC and IMC are derived. Here, the approximated IMC control signal $\hat{y}'_a(n)$ is compensated on the residual error signal $e(n)$ before it is fed into the MVC controller. This isolates the MVC from the influence of the IMC. Simultaneously, the input to the IMC controller, $\hat{d}(n)$, is not an approximation of $d(n)$, but instead the approximation of the remaining residual error of the MVC control. It is clear that the two additions on the bottom generate a duplicate, but this is justified after seeing the mixed digital-analogue implementation presented in Fig. 2(a), where one is calculated and used in the digital domain and the other in the analogue one.

The concatenated control of the IMC over the MVC residual noise combines the MVC transfer function and its own transfer function into a new one that consist of the multiplication of both. This concatenates the attenuation (and amplification) that each controller would produce by its own means, providing at the same time that the controllers can be now designed independently from each other, allowing a very uncomplicated process of on-line optimization and adaption of both control systems.

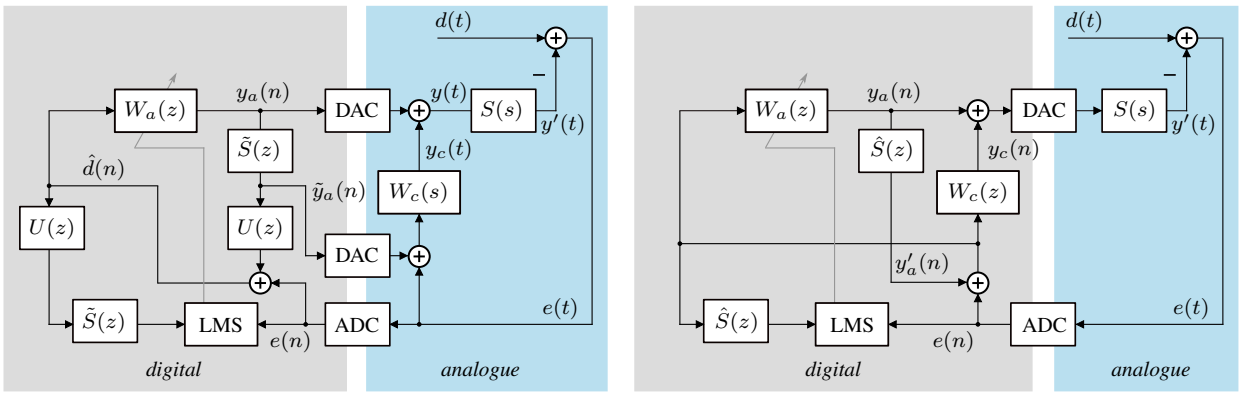

Figure 2: System diagrams of (a) the mixed analogue-digital implementation [7] and (b) the fully digital implementation.

The first implementation of this system is depicted in Fig. 2(a). It comprehends an analogue 2nd order fixed MVC implementation combined with a digital adaptive FIR IMC implementation based on an FxLMS. A hybrid implementation was chosen because of the negative effect that the processing and ADC and DAC conversion delays would have on the attenuation bandwidth of the MVC controller. This requires the DAC conversion of the filtered control signal $\tilde{y}_a(n)$ to generate the input to the MVC controller in the analogue domain. For this, a separate measurement of the secondary path $S(z)$ and the effect of the ADC and DAC conversion are needed, respectively denominated in the diagrams as $\tilde{S}(z)$ and $U(z)$. In the fully digital implementation depicted in Fig. 2(b), there is no need for an additional DAC and the independent approximations of $S(z)$ and $U(z)$. Instead, a single measurement $\hat{S}(z)$ is needed to represent the effect of the secondary path and the conversions together. Moreover, a duplicate of the controllers' input calculation is not longer required and the two filters storing $U(z)$ can be neglected, since all happens within the digital domain.

All in all, by implementing the MVC in the digital domain the complication of having to perform two different system measurements, instead of only one, and the additional cost of a second DAC unit can be avoided. At the same time, the analogue complexity of the overall structure decreases significantly and offers also the opportunity to fine-tune the MVC without having to solder a new circuit. All these reasons motivate the present effort to find a faster ADC and DAC conversion technology and signal processing techniques to implement a fully digital system.

3. Towards Low-Latency Processing

The discrete-time filtering of a causal recursive filter and its input signal $x(n)$ can be written as two separated sums

$$y(n) = \sum_{k=0}^{K} b_k \cdot x(n-k) - \sum_{l=1}^{L} a_l \cdot y(n-l), \tag{1}$$

where the first one describes the interaction of its input samples with its feed-forward filter coefficients b_k, and the second one the interaction of its past outputs with its feedback coefficients a_l. For the case of a non-recursive filter, e.g. FIR filter, the second sum may be ignored. Depending on the order of the filter, the computational power of the target platform, and the strategy of calculation, the time needed to calculate both sums of products for one input sample, called the processing latency, may increase beyond the sampling period or decrease below it. The first case would produce that the result of the convolution will not be ready before the new input sample comes, while the second case would produce a result before the next sample arrives. If ADC and DAC units are clocked with the same signal, i.e. coupled together, the processing latency will have an *effective* value that is a

multiple of the sampling period and greater or equal to its *real* value. To get rid of this undesirable extra processing latency and reach an *effective* value equal to the *real* value, the ADC and DAC units should be decoupled.

If the sampling period is as long as or longer than the time required to calculate the convolution, then the input and output samples involved in the convolution do not change during its calculation. In that case, the first term of the left sum of products may be pulled out

$$y(n) = b_0 \cdot x(n) + \sum_{k=1}^{K} b_k \cdot x(n-k) - \sum_{l=1}^{L} a_l \cdot y(n-l) \tag{2}$$

to show that besides $x(n)$, all other input samples were already known before starting with the calculation. If this fact is strategically used to provide the next convolution calculation with a precalculated output

$$\tilde{y}(n+1) = \sum_{k=1}^{K} b_k \cdot x(n-k+1) - \sum_{l=1}^{L} a_l \cdot y(n-l+1), \tag{3}$$

taking care of including the current output $y(n)$ into the right sum, then the processing latency is minimized to the time it takes to perform one multiplication and one addition

$$y(n+1) = b_0 \cdot x(n+1) + \tilde{y}(n+1) \tag{4}$$

without having to allocate additional computational resources.

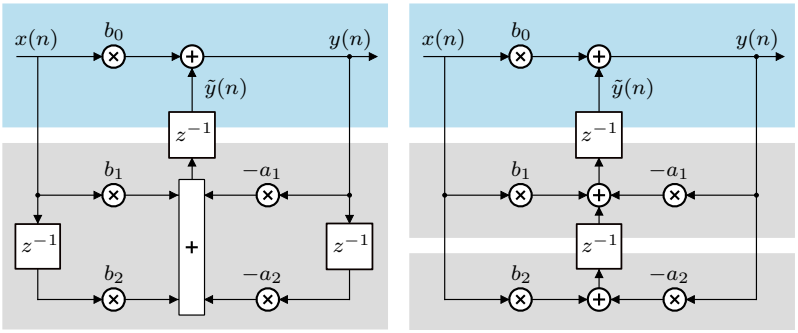

Figure 3: Special cases of a 2nd order IIR filter: (a) Hybrid Form I and and (b) Transposed Direct Form II.

If this simple strategy is implemented with an IIR Direct Form I subblock for the calculation of $\tilde{y}(n+1)$, then the implementation is a special case of an IIR Hybrid Form I [9] (see Fig. 3(a)). In this case we will say that the convolution calculation is *decoupled* from the filter's memory. If the precalculation is applied iteratively, then the behavior matches the well-known IIR Transposed Direct Form II (see Fig. 3(b)).

4. MVC Digital Implementation

The platform used for the implementation is the MicroLabBox from dSpace. The platform consists of a Xilinx Kintex-7 FPGA connected through a local bus to a Freescale dual-core 2 GHz processor that makes the interface for controlling and programming the platform from an external PC. Both, processor and FPGA, are graphically programmed using Simulink/System Generator. The FPGA runs with an internal clock rate of 100 MHz and has direct access and control over the ADC and DAC units. Both are built with SAR technology and provide delays of only 1 μs [10]. They operate with 16 bit quantizers and an analogue range of ± 10 V. The digital signals are represented in two's

complement and a fixed-point representation of *Fix*16_0, where the *Fix* classifies a signed fixed-point representation and the two numbers separated by an underscore denote the total number of bits and the number of decimal bits respectively. Therefore, the maximum input of 10 V is mapped to a digital value of $2^{15}-1 = 32\,767$. The ADC unit is programmed through the FPGA with a 48 008 Hz periodic signal to sample and deliver the error signal $e(n)$, and the DAC unit is connected to the filter structure to convert its output when it is ready.

4.1 Control Logic

The Moore State Machine in Fig. 4 controls the internal signal processing of the IIR filter structure shown in Fig. 3(a).

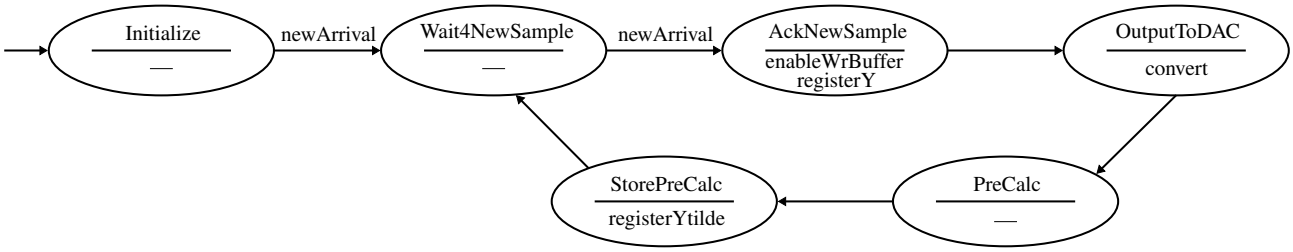

Figure 4: Moore State Machine of the IIR-Filter

In state *Initialize* the machine waits for the initialization of the ADC. When this is done the machine gets a *newArrival* signal and switches to state *Wait4NewSample*. In this state the machine remains until a *newArrival* signal announces the new sample. Due to a falling edge detector, which is needed to limit the *newArrival* signal of the platform to one clock cycle, the *newArrival* signal is delayed by one clock cycle in comparison to the new sample. The filter is implemented in a way, such that the calculation of the current output (see Eq. (4)) is done all the time. Therefore, the calculation with the new sample starts one clock cycle before the corresponding *newArrival* signal arrives, while the state machine is still in the *Wait4NewSample* state. When jumping into the state *AckNewSample* the current output has already been calculated and can be stored into the output register. In addition, this state updates the buffer of the previous in- and outputs. After one clock cycle the state machine changes to state *OutputToDAC*. In this state the registered output and a *convert* signal are transmitted by the filter. The *convert* signal carries the information that the current filter output is available. Therefore, this signal is used by the DAC as starting signal for the conversion. With this implementation simply three clock cycles are needed to produce the current filter output, thus the filter yields a signal processing latency of only 30 ns.

After outputting the result the state machine jumps into state *PreCalc* for seven clock cycles. This time is needed to precalculate $\tilde{y}(n+1)$ for the next sample according to Eq. (3). When the precalculation is done, the state machine changes to state *StorePreCalc* and stores the precalculated value in a register. In the next clock cycle the state machine ends up in state *Wait4NewSample*, where the filter waits for a the next sample.

Between the IIR filter and the DAC the filter output is multiplied by a negative factor. The minus sign is needed to produce the anti-phase for the noise cancellation process and the magnitude can be used as an additional digital gain.

4.2 IIR filter and effects of quantization

To attenuate the low frequencies a 2$^{\text{nd}}$ order IIR filter is used for the MVC. The complex poles and zeros are designed with frequencies of $f_p = 700$ Hz and $f_z = 1$ kHz. Their magnitudes are $r_p = 0.99286$ and $r_z = 0.961985$, respectively. This setting yields the magnitude and phase responses

shown in Fig. 5(a) and Fig. 5(c). The corresponding coefficients can be calculated to

$$b_0 = 0.308674367063210\ ,\ b_1 = -0.588799494372055\ ,\ b_2 = 0.285651932679664 \tag{5}$$

and

$$a_1 = 1.97738976078356\ ,\ a_2 = -0.985770979600000. \tag{6}$$

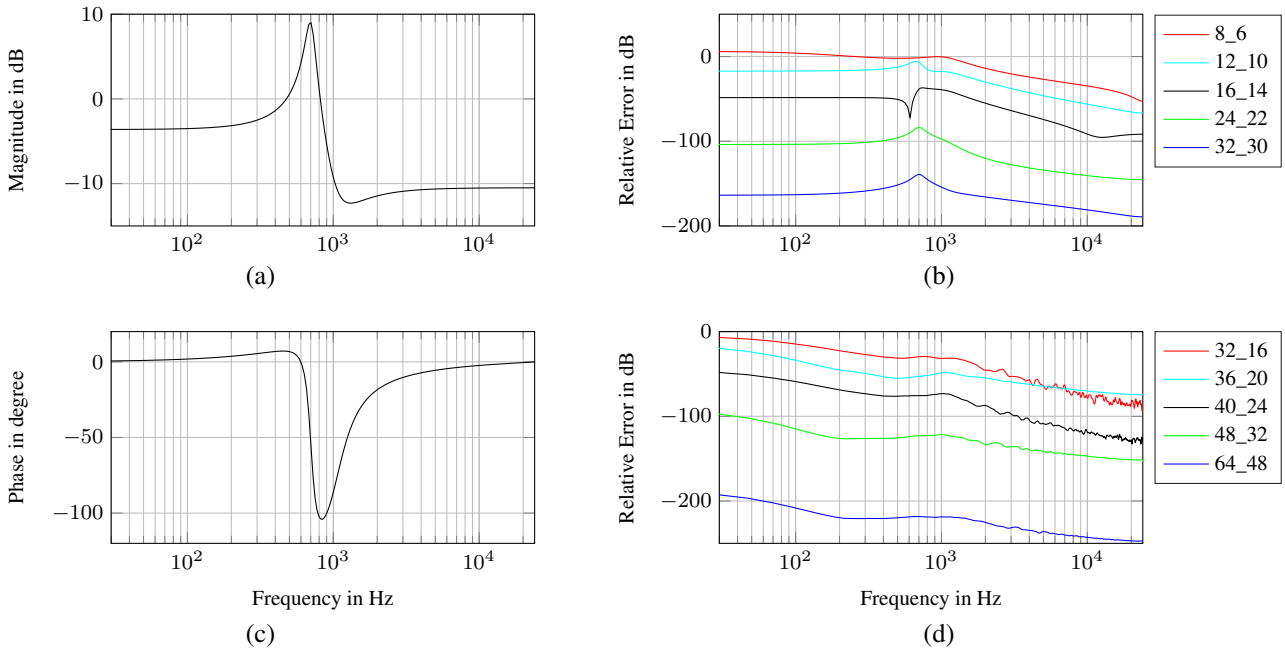

Figure 5: (a) Magnitude and (c) phase response of the theoretical IIR filter, together with (b) relative error of the frequency response for different coefficients' quantizations, and (d) relative error of the frequency response for different quantizations of the output

In order to be used on the FPGA, the coefficients have to be represented in fixed-point representation. Thus the coefficients have to be quantized to a chosen representation. At the same time, it is important to notice that every quantization of the coefficients results in a change of the poles and zeros. This change in the parameters leads to changing properties of the filter. For this reason different quantizations are simulated in MATLAB. Fig. 5(b) shows the relative error in the magnitude response for different quantizations. It can be seen that an increasing bit-width of the coefficients decreases the relative error but at the same moment increases the computational complexity. Hence, a reasonable trade-off between error in magnitude response and computational complexity has to be found. The threshold for the maximum relative error has been chosen to -60 dB. The curve which fits this criteria, is the one with a representation of *Fix*24_22, wherefore this quantization is chosen for the coefficients.

A second quantization has to be done before the output sample reaches the feedback loop. Otherwise, the bit-width of the filter output will increase after every iteration of the feedback loop. To simulate this quantization, the coefficients are represented in the chosen *Fix*24_22 representation. The resulting errors in magnitude response for different quantizations are shown in Fig. 5(d). To achieve the same threshold of at most -60 dB relative error for every frequency a quantization of *Fix*48_32 has been selected.

5. Measurement Results and Evaluation

The MVC implementation with the coefficients presented in Section 4 and an adaptive FxLMS IMC of length 4096, which can not be described here due to space restrictions, are loaded on the MicroLabBox platform. To enable the IMC to work, a non-recursive internal model $\hat{S}(z)$ of length

8192 taps is measured and derived with the technique and headphone prototype described in [11]. Similarly, the prototype is placed on a Neumann KU100 dummy-head and connected to the platform. To deliver the current needed to drive the headphones' speakers, a Behringer Powerplay Pro-8 power amplifier is connected between the MicroLabBox's output and the jack of the headphones. A Genelec 8030B loudspeaker is placed facing towards the left ear-cup, 70 cm away from the dummy-head inside of a room designed for audio-listening with dimensions 4.80 x 4.20 x 2.05 m. Uniformly distributed pseudo random noise with 20 dB louder sinusoids of frequencies 105.5, 398.5, and 2004.2 Hz are played through the speaker, generating 89 dB_{SPL} at the location of the ear-cup.

The system is additionally programmed to deliver 48 times per second a 4096 taps long FFT of the measured residual error signal $e(n)$ for evaluation purposes (see Fig. 2(b)) and to be switched between the following states: *OFF*, no output sample is given to the DAC to be transformed; *MVC*, only the MVC is working and the signals $y_a(n)$ and $y'_a(n)$ coming from the IMC substructure are being ignored in the calculations; *IMC*, only the IMC is working, the signal $y_c(n)$ is ignored in the calculation, and the adaption of the filter can be turned on, off or reseted; and *HYBRID*, where both substructures are working together.

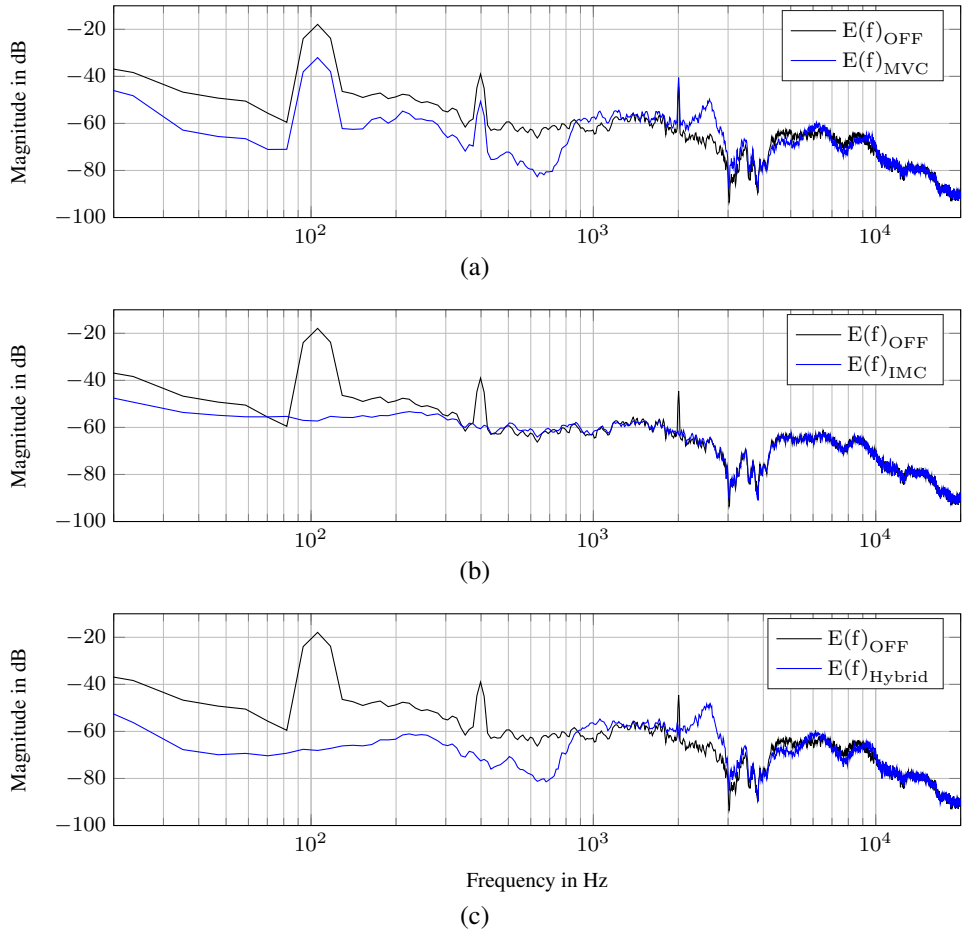

Figure 6: Error signal's measured spectra under (a) MVC control, (b) IMC control, and (c) hybrid control. Disturbance noise is uniformly distributed pseudo random noise with 20 dB louder sinusoids of frequencies 105.5, 398.5, and 2004.2 Hz.

The measurement results are presented in Fig. 6. The plots are 2048 taps long single-sided FFT spectra of the error signal over the 20-20 kHz frequency range represented in dB. The three different control states (*MVC*, *IMC*, and *HYBRID*) are compared with $e(n)$ when the system is in state *OFF*, which is equivalent to $d(n)$, if the noise conditions remain constant. In Fig. 6(a) it can be seen that the MVC controller produces a relative constant attenuation between 10 and 20 dB over the low frequency range up to approximately 870 Hz, but producing an amplification of 10 dB around the

2600 Hz, which is a known behaviour of MVC controllers. In Fig. 6(b) the attenuation generated by the IMC after 15 minutes of adaption can be seen. Its effectiveness on periodic signals is very present with 40 dB, 21 dB, and 22 dB attenuation for the respective frequencies 105.5, 398.5, and 2004.2 Hz. A non-uniform broadband attenuation with a peak of 10 dB within the range 20-340 Hz is also observed. In Fig. 6(c) the attenuation of the hybrid structure can be seen. As expected, the attenuation (and amplification) of both individual control substructures are combined, reaching constantly values between 10 and 50 dB. Nevertheless, deviations around the perfect combination of the individual substructures are also found. This deviation has a mean value of -0.01 dB and standard deviation of 1.34 dB, which is most probably generated by the internal fixed-point number representation of the signals and calculations, and the fact that the measurements could not be done simultaneously.

All in all, the fast ADC and DAC units, their decoupled control, the convolution calculation strategy, and its implementation on an FPGA platform decreased the overall latency to 2.03 μs, consisting of 1 μs for the ADC, 30 ns for the filtering process, and another 1 μs for the DAC. This latency was validated to be small enough to produce comparable results to an analogue implementation of the MVC, with an attenuation bandwidth of 870 Hz and values between 10 and 20 dB. At the same time, the fully digital implementation decreased the need for multiple system measurements and an extra DAC unit.

REFERENCES

1. Rafaely, B. Active noise reducing headset - an overview, *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, **2001** (3), 2144–2153, (2001).

2. Pawełczyk, M. Analogue active noise control, *Applied Acoustics*, **63** (11), 1193–1213, (2002).

3. Pawełczyk, M., Elliott, S. and Rafaely, B., Active noise control using feedback. Fixed and adaptive controllers, (1997).

4. Elliot, S. J., *Signal Processing for Active Control*, Signal Processing and its Applications, Academic Press, London (2001).

5. Pawełczyk, M. A hybrid active noise control system, *Archives of Control Sciences*, **13** (2), 191–213, (2003).

6. Song, Y., Gong, Y. and Kuo, S. M. A robust hybrid feedback active noise cancellation headset, *IEEE Transactions on Speech and Audio Processing*, **13** (4), 607–617, (2005).

7. Schumacher, T., Krüger, H., Jeub, M., Vary, P. and Beaugeant, C. Active noise control in headsets: A new approach for broadband feedback ANC, *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May, pp. 417–420, (2011).

8. Liebich, S., Anemüller, C., Vary, P., Jax, P., Rüschen, D. and Leonhardt, S. Active noise cancellation in headphones by digital robust feedback control, *Proceedings of the 24th European Signal Processing Conference (EUSIPCO)*, Sept, (2016).

9. Azadet, K. and Nicole, C. J. Low-power equalizer architectures for high speed modems, *IEEE Communications Magazine*, **36** (10), 118–126, (1998).

10. dSpace GmbH, (2015), *MicroLabBox Features*. Release 2015-A.

11. Rivera Benois, P., Bhattacharya, P. and Zölzer, U. Derivation technique for headphone transfer functions based on sine sweeps and least squares minimization, *Proceedings of the 45th International Congress and Exposition on Noise Control Engineering, INTER-NOISE*, Aug, (2016).