Implementation of a New Method to Digital Audio Equalization

Gabriel F. P. Fernandes*,
Luís Gustavo P. M. Martins*,
Miguel F. M. Sousa*,
Filipe S. Pinto*,
Aníbal J. S. Ferreira[o]*,

[o]Department of Electrical and Computer Engineering of the University of Porto
*INESC Porto, Largo Mompilher 22, 4000 Porto, PORTUGAL

Abstract

This paper describes the design and implementation of an FIR based 20-band stereo audio equalizer that offers good spectral selectivity and phase response. The underlying theory of fast filtering in the frequency domain is addressed, the architecture of the DSP system using a single TI TMS320C31 processor is presented, the functionality of the interface, either software-based or hardware-based is described, and a few test results are presented and discussed.

1 Introduction

Digital audio equalizers offer many advantages over analog equalizers [1][2]. In fact, in terms of quality, typical problems of analog equalizers like group delay distortion, interband interference, difficulty to achieve a perfectly flat frequency response, poor global system controllability, may be completely eliminated or attenuated within some desired mathematical precision by using digital signal processing techniques. On the other hand, it is possible to design very efficient high quality digital filters using multirate structures [3] or by using well known techniques of fast filtering in the frequency domain [4]. For efficiency reasons, low-order Infinite Impulse Response (IIR) filters are frequently used to implement a desired equalization curve. However, these filters have a rather poor phase response, and generate disturbing signal discontinuities when their coefficients are switched abruptly [9].

In this paper we will consider Symmetric Finite Impulse Response (FIR) filters for digital equalization because they do not suffer from this latter problem and because they are inherently linear phase and thus they eliminate group delay distortions.

Several digital equalizers use the coefficients of the impulse response of the FIR equalization filter to perform a direct linear convolution in the time domain with the audio data. The number of bands of this type of equalizers is normally low (e.g. 10) because very high spectral selectivity implies very long digital filters which means a proportionally higher computational load and longer system delays that may not suit the intended application.

A much more efficient implementation of linear convolution may be achieved by using frequency domain techniques of fast filtering like the known methods of

overlap-and-add (OLA) and *overlap-and-save* (OLS) [4]. However, several precautions must be taken to insure that the circular convolution that is in fact implemented by these methods corresponds in practice to the linear convolution. This implies that after a modification of the frequency response settings of the equalizer, a new filter must be designed [5]. This step involves normally a considerable computational effort, and the final equalization filter is always an approximation to the desired equalization settings.

This paper addresses the real-time implementation of a modified method of fast filtering in the frequency domain. This method simplifies the computation of the discrete coefficients of the frequency response of the equalization filter from the specified equalization curve. The obvious benefit is that the equalization filter is easily computed by the DSP that is charged of the overall processing. The obvious consequence is that aliased components should be expected. However, instead of seeking a mathematical elimination of these components, as the classic methods of OLA or OLS do, our modified method of fast filtering seeks to eliminate its perceived presence in the processed signal [6].The block diagram of Figure 1 illustrates this approach. It has been used to design a *20-Band Digital Audio Equalizer* as described in [7].

This paper is structured as follows: Section 2 addresses the system architecture and describes the implementation and functionality of the DSP system and user interface. Section 3 presents the results of a few tests that illustrate the performance of the equalizer and Section 4 explains our plans to extend the system to more sophisticated processing. Section 5 concludes this paper.

2 System Architecture

The real-time implementation of an FIR based *20-Band Digital Audio Equalizer* implies the consideration of a number of demanding requirements in order to accomplish a glitchless and stereo high-quality audio (24 bit, 44.1 kHz) processing. For example, the calculation of the equalization filter coefficients (based on the user defined equalization curve), as well as the corresponding filtering of the audio signal, must be real-time operations. Furthermore, the audio input/output operations must be managed flawlessly in order to preserve the high-quality audio stream.

The real-time implementation of the *20-Band Digital Audio Equalizer* is based on a single low cost Texas Instruments TMS320C31 32 bit floating point digital signal processor running at 50 MHz (30 MIPS, 60 MFLOPS). The choice of a 32 bit floating point DSP has been motivated by its ability to handle signals of differing precision, dynamic range, and signal-to-noise ratios, and also for simplifying the development phase. The system architecture is depicted in Figure 2. It consists of high-quality A/D and D/A stereo converters operating at 44.1 kHz sampling frequency, 64 Kword FLASH RAM, 128 Kword fast RAM and a parallel port logic controlling the communication of the processing system with the user interface. This communication is bidirectional and involves the transfer of equalization coefficients (from the user interface to the processing system) and spectral energies (from the processing system to the user interface).

This system is similar to the one the authors have submitted to the "*1997 Texas Instruments DSP Solutions Challenge*" – "*An Efficient 20-Band Digital Audio Equalizer"*, which has been selected as one of the European "*Top 20 Projects*" [7].

A/D and D/A Conversion

A/D and D/A 24 bit converters are used in the system (Crystal CS5360 and CS4390, respectively). It is also possible to use 20 bit A/D and D/A converters (quality/cost compromise) since Crystal Semiconductors provides pin compatible 20 bit A/D and D/A converters.

The A/D and D/A converters are connected without glue logic to the 'C31 serial port as illustrated in Figure 3.

The CS5360 A/D converter is set to MASTER mode, generating internally the word and bit clock signals (LRCK and SCLK respectively) from the MCLK signal and supplying it to the entire system. The 'C31 serial port and the CS4390 D/A converter are configured as SLAVE devices and expect to receive the CS5360 generated word and bit clock signals at their respective inputs (unlike the 'C31 serial port, the CS4390 can only work as a SLAVE device).

The 'C31 receives the digital audio data from the A/D, processes it, and outputs the processed audio to the D/A.

S/PDIF AES/EBU Digital Audio Interface

It is also possible to connect an **S/PDIF- AES/EBU Digital Audio Interface** to the system. This solution provides a noiseless audio interface to the outer world (e.g. CD Players and DAT recorders with S/PDIF-AES/EBU outputs/inputs) since the noise introduced by A/D and D/A converters can be avoided: the audio remains in the digital domain during the whole path.

Although an S/PDIF interface has already been designed and successfully tested in our system, in the current implementation it is not possible to simultaneously have the A/D and D/A converters and the S/PDIF Digital Audio Interface connected to the system since the 'C31 provides only one serial port. In a future version of the system it will be possible to have simultaneous analog and digital inputs/outputs (the user will be able to choose which audio input to use: the analog or the digital one).

Parallel Port Interface

A parallel port interfaces the system allowing a bidireccional connection to a PC where a software interface is running, or to a hardware-based user interface.

Each equalizer audio channel (left/right) has its own independent 20 band equalization curve. This results in 40 values that are transmitted from the user interface, either hardware or software based, to the DSP memory (through the parallel

port) in each communication burst. These 40 values are always transmitted when a modification occurs on any of the equalization curves of either channel.

A *Spectrogram* view is available in the user interface of the *20-Band Digital Audio Equalizer*, where the DSP calculated energies in each band are displayed in real-time. In order to accomplish this, 40 values corresponding to the left and right channels (20 values for each channel) are sent to the User Interface every 70 ms.

External RAM and FLASH Memory

External zero-wait-state RAM memory and FLASH memory are also an important part of the system. The existence of external RAM is mandatory since the ‘C31 2Kwords of internal memory is quite insufficient to accommodate any complex audio processing code. Zero-wait-state memories are used in order to attain real-time operation. FLASH memories allow a greater flexibility over the more common EPROMs, because they allow boot-up code upgrade by software, thus avoiding hardware intervention.

Optimizing Critical Code

The *20-Band Digital Audio Equalizer* uses computationally intensive techniques such as the FFT algorithm. We were compelled to optimize the code in order to be able to process stereo digital audio in real-time at 44.1 kHz (in blocks of 512 new samples, *i.e.* every 11.6 ms). A clear understanding of the TMS320C31 DSP run-time restrictions and capabilities was therefore mandatory.

Routines to calculate the ODFT (Odd-Discrete Fourier transform) and the IODFT (Inverse Odd-Discrete Fourier Transform) [6] ended up to be implemented in the internal memory of the DSP, for faster execution. The manual optimization of the assembly code regarding critical processing became also mandatory.

Other important aspects concern pipeline conflicts. Correct handling of branch, register and memory conflicts allow a better algorithm performance. Before the ODFT computation, data buffers to be processed are copied into internal memory, since write instructions to external memory require two bus cycles to complete, instead of one.

The use of Decrement and Branch Delayed instructions (DBD) in replacement of nested Repeat Block (RPTB) ones (these require multiple PUSH and POP instructions for context switching) also led to time savings.

Along the code, several manipulations of data buffers take place. These are reduced to a minimum, since we deal with large buffers whose size is 1024 samples. These copies consume considerable amount of time, and for efficiency they are either included in other operations performed to the buffers, or implemented using DMA operations.

Particular attention was given to the parallel execution of the maximum possible number of instructions. The butterfly loop of the ODFT/IODFT is the most intensive

part of the algorithm. Having this particular part implemented with as many parallel instructions as possible, resulted in a great performance improvement.

Audio Samples IN/OUT

Aiming at CD Audio quality, samples are converted at a sampling frequency of 44.1 kHz. The size of each processing audio frame is 512 samples, which forces all the processing to be done in less than 11.6 ms.

The same interrupt service routine (ISR) responsible for receiving a new 24 bit audio sample, is also responsible for sending the last processed one and ready for output. This solution is preferred over the use of separate ISR for receiving and for transmitting audio samples, allowing faster performance, since the pipeline needs to be flushed and the context switching needs to be performed every time an ISR is serviced. In addition, extra time saving is also achieved by placing the ISR in internal memory.

Hardware and Software User Interface

A great deal of effort has been placed on the design and functionality of the user interface. In order to improve the flexibility of the system, two interfaces have been implemented, one being software-based (Win95/98), the other being hardware-based. Both interfaces have the same look and allow the same functionalities and real-time operation. Figure 4 illustrates the design of these interfaces.

Both interfaces allow the user to edit independent equalization curves for each audio channel (left/right) the way he would do on an analog device: each band has a slider that determines its gain/attenuation value in the range [-12, +12] dB. These sliders can be adjusted on the fly (by using the buttons and the dial on the hardware interface, or by clicking and dragging the sliders on the software version) and the corresponding modifications are reflected instantly in the processed audio, without any glitches or other switching artifacts. When scrolling a band's slider the corresponding gain/attenuation value and central frequency information are displayed on a pop-up window for precise adjustment.

One flexible mode of adjusting equalization curves is the *Stereo-Link Mode*. When in this mode, any modifications made to one of the audio channels (left/right) will be reflected to the other channel, even if they have different curves. The modification passed to the other channel is not absolute, being relative to its own actual value. In other words, only the decrements/increments are passed to the linked channel, thus not destroying its equalization curve.

The equalization curves of both channels can be stored and latter recalled from memory. The memory menu provides a convenient way of selecting the stored curves by means of a thumbnail preview of each equalization curve. It also allows the user to activate temporarily any stored curve, without loosing the actual one, before loading it. Up to 12 stereo equalization curves can be stored in the system.

The interface has more features that include the possibility to swap equalization curves between channels and to copy curves from one channel to the other.

Another mode of operation is the Spectrogram view. Spectral power values of each band of each channel are calculated and displayed in real-time by the DSP. It is possible to view simultaneously the Spectrogram of both channels or of each channel at a time for a more detailed view.

In all of the Spectrogram view modes it is possible to watch real-time information about each band of each channel (spectral power and central frequency), displayed in a floating box, by using the buttons on the hardware interface or by clicking in the desired band on the software interface.

The equalization curves and the Spectrogram values on the user interface are transmitted/received to/from the DSP through the PC parallel port connection, providing a fast and reliable communication between these devices.

3 Evaluation Tests

The evaluation tests were performed considering different equalization curves and by comparing the objective performance of our *20-Band Digital Audio Equalizer* against that of a *31-Band Digital DSP Equalizer* which is considered a market reference and is manufactured by a well-known and famous audio equipment manufacturer. For the sake of simplicity this system will be, from now on, referred to as *Equalizer X*.

The measurements were done using a *Precision Audio System One Plus* Audio Analyzer (SNR, THD+N and cross-talk measures) and a *SRS SR780* Network Signal Analyzer (frequency response and group delay plots).

The tests results are presented in Table 1 and the frequency response and group delay plots are depicted in Figures 5, 6 ,7 and 8.

Referring to Table 1, the *20-Band Digital Audio Equalizer* presents an SNR value of 108dB (@1kHz, 10dBu), an THD+N maximum value of 0.004% (@1kHz , 10dBu) and a cross-talk figure of 98dB. These values indicate that with respect to the SNR and THD+N criteria, the *20-Band Digital Audio Equalizer* and the *Equalizer X* have a comparable performance. With respect to the cross-talk the former performs better than the latter.

The frequency response and group delay plots depicted on Figures 5 and 6 represent the response of each system to a flat equalization curve (all equalizer bands set to 0 dB attenuation/gain value). In this case no group delay distortions are measurable on both devices and their frequency response is flat. Their frequency bandwidth is similar and goes up to 22kHz (0dB ± 0.5dB).

Figures 7 and 8 depict the frequency responses and group delay plots of the *20-Band Digital Audio Equalizer* and of the *Equalizer X*, when the equalization curve

presented in Table 2 is set on both devices. This equalization curve is the same one presented and used as an example in [6].

The differences between the responses of the two systems tested are now quite evident.

First of all, the frequency response of *Equalizer X* and regarding the first three bands is quite puzzling. In fact, although its frequency selectivity is superior to that of the *20-Band Digital Audio Equalizer* for the third band, it appears to not being able to set the first band (centered at about 50 Hz) to +8dB. On the other hand, the *Equalizer X* exhibits low precision, as well as interband interference (Figure 8). As example, the desired gain value for the 8^{th} band (centered around 1 kHz) should be 12 dB, when the Table 2 equalization curve is set on the user interface. This value is not achieved in practice as illustrated in Figure 8, where the *Equalizer X* has a 10.027 dB instead of 12 dB gain at 1.024 KHz. This results in an equalization curve somewhat different from the desired one. Looking to the group delay plot the non-linear phase response for the *Equalizer X* becomes evident, which is a typical characteristic of IIR filtering techniques. This results in group delay distortions with negative impact on the audio quality.

On the other hand, the FIR based *20-Band Digital Audio Equalizer* presents a very well defined and selective frequency response (Figure 7) that results in a truthful equalization curve, taking as a reference the one set by the user. The group delay (Figure 7) remains constant whatever equalization curve is set introducing no group delay distortions to the audio signal.

These results confirm that FIR based filters exhibit a much better performance than IIR filtering techniques, particularly in what concerns phase response linearity and the corresponding group delay distortion. This is important in quality sensitive applications.

4 Future Developments

Future work will be directed to extend the system to more sophisticated processing such as automatic room equalization. This automated equalization system will include both magnitude and phase compensation, minimizing phase errors in the loudspeakers and crossovers, as well as allowing automated magnitude response equalization based on the room acoustic response.

Parametric equalization can also be added to the system, making it easy the inclusion of an automatic feedback exterminator, which is very useful in live performance situations.

Also planed are some improvements to the User Interface:

- simultaneous modification by the same amount of a user defined group of bands for a faster and more convenient curve editing;
- an automatic clipping reduction function that will avoid distortions produced by excessive gain settings.

5 Conclusion

We have described in this paper the implementation on a single low cost TMS320C31 DSP of a 20-Band Digital Audio Equalizer that is based on a modified technique of fast filtering in the frequency domain.

To achieve real-time operation, manual critical code optimization and low-level assembly language programming has been necessary.

The use of top notch A/D and D/A converters is also necessary to achieve high-quality audio specifications comparable to those of equipments that are often pointed as market references in the area of high-quality digital audio equalization. The expected good spectral selectivity and phase response of an FIR based system like the one presented in the paper were absolutely confirmed by the performance test results carried out and described in this paper.

An intuitive and powerful user interface has also been described that improves the system flexibility.

Future functionalities will be integrated and will provide extra-value to the *20-Band Digital Audio Equalizer*. Even if it is in its very first version this system already proves to be a low-cost, user-friendly, high-quality digital audio processing system, which may compete with higher-cost equipments that perform similarly.

References

[1] Terry E. Riemer Juan A. Henriquez and Russel E. Trahan. “A Phase-Linear Audio Equalizer: Design and Implementation”. Journal of the Audio Engineering Society, 38(9):653-666, September 1990.

[2] Paul H. Kraght. “A Linear-Phase Digital Equalizer with Cubic-Spline Frequency Response”. Journal of the Audio Engineering Society, 40(5):403-414, May 1992.

[3] Ronald E. Crochiere and Lawrence R. Rabiner. “Multirate Digital Signal Processing”. Prentice-Hall, 1983.

[4] Alan V. Oppenheim and Ronald W. Schafer. “Digital Signal Processing”. Prentice-Hall Inc., 1975.

[5] James H. McClellan, Thomas W. Parks, and Lawrence R. Rabiner. “A Computer Program for Designing Optimum FIR Linear Phase Digital Filters”. Transactions on Audio and Electroacoustics, AU-21(6):506-526, November 1973.

[6] Aníbal J. S. Ferreira and José M. N. Vieira. “An Efficient 20-Band Digital Audio Equalizer”. Presented at the 98th AES Convention, Preprint 3992, February 1995.

[7] INESC Porto – Portugal. “An Efficient 20-Band Digital Audio Equalizer”. 1997 DSP Challenge, European Top 20 Entries, Texas Instruments, 1998.

[8] “TMS320C3x User’s Guide”. Digital Signal Processing Products, Texas Instruments, 1994.

[9] Vesa Valimaki and Timo Laakso. “Supression of Transients in Time-Varying Recursive Filters for Audio Signals”. IEEE, 1998.

Table 1: Specifications.

	20-Band Digital Audio Equalizer	*Equalizer X*
SNR (@1kHz, 10dBu)	108 dB	108 dB
THD+N (@1kHz, 10dBu)	< 0.004 %	< 0.004 %
CROSS-TALK	98 dB	80 dB

Table 2: Center frequency of each band-pass filter and an example of equalization settings.

BAND	CENTER (Hz)	SETTING (dB)
1	47	8
2	140	0
3	234	-8
4	352	0
5	492	0
6	633	-12
7	820	0
8	1008	12
9	1242	0
10	1523	-2
11	1875	-8
12	2320	-4
13	2930	5
14	3680	0
15	4805	0
16	6352	0
17	8507	0
18	11414	0
19	15210	0
20	20883	0

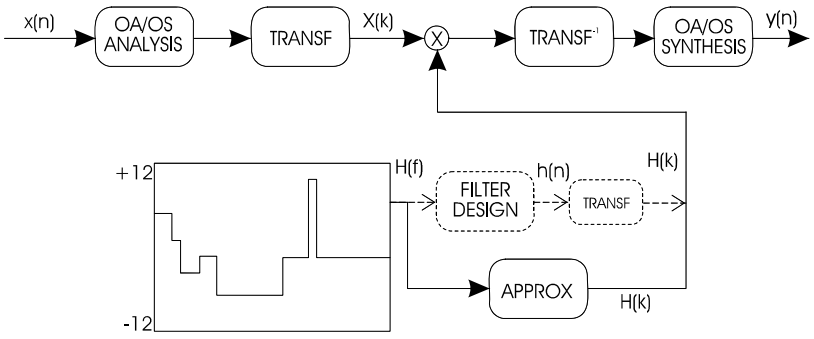

Figure 1: Algorithm structure.

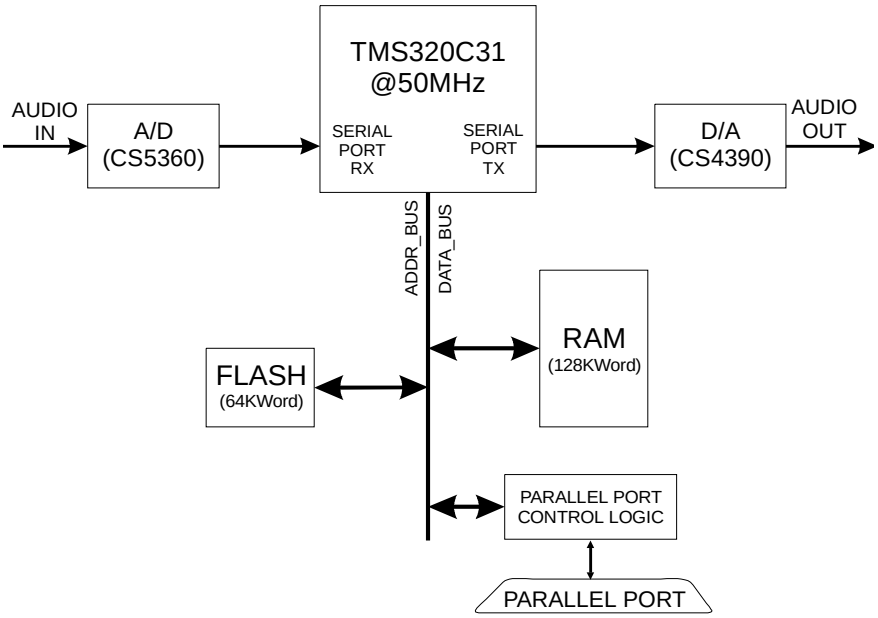

Figure 2: System architecture.

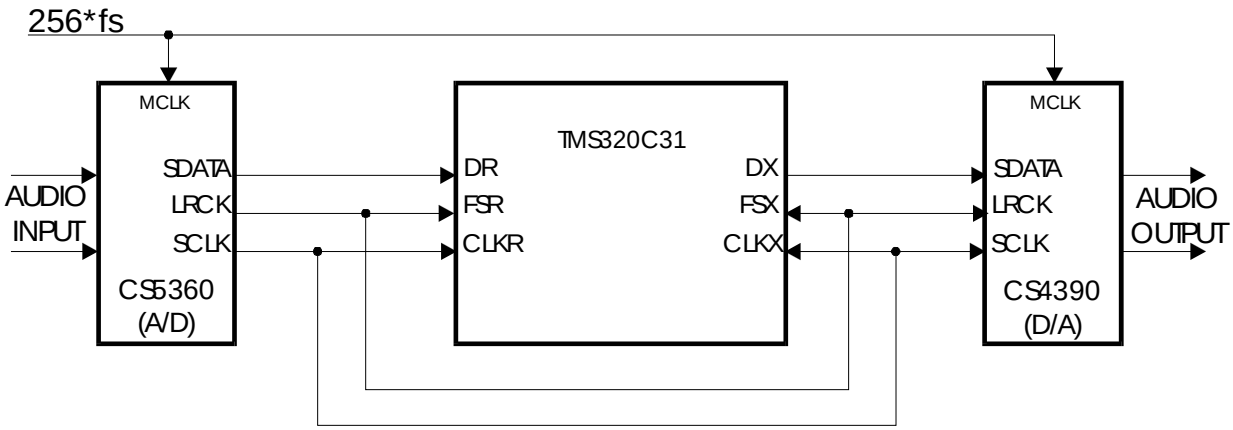

Figure 3: interfacing the 'C31 with the A/D and D/A converters.

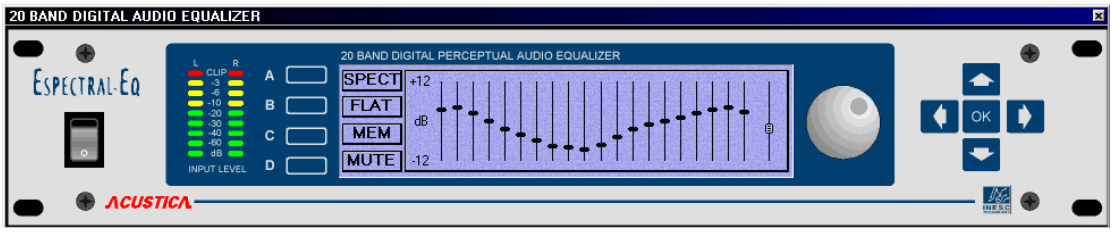

Figure 4: Graphical User Interface.

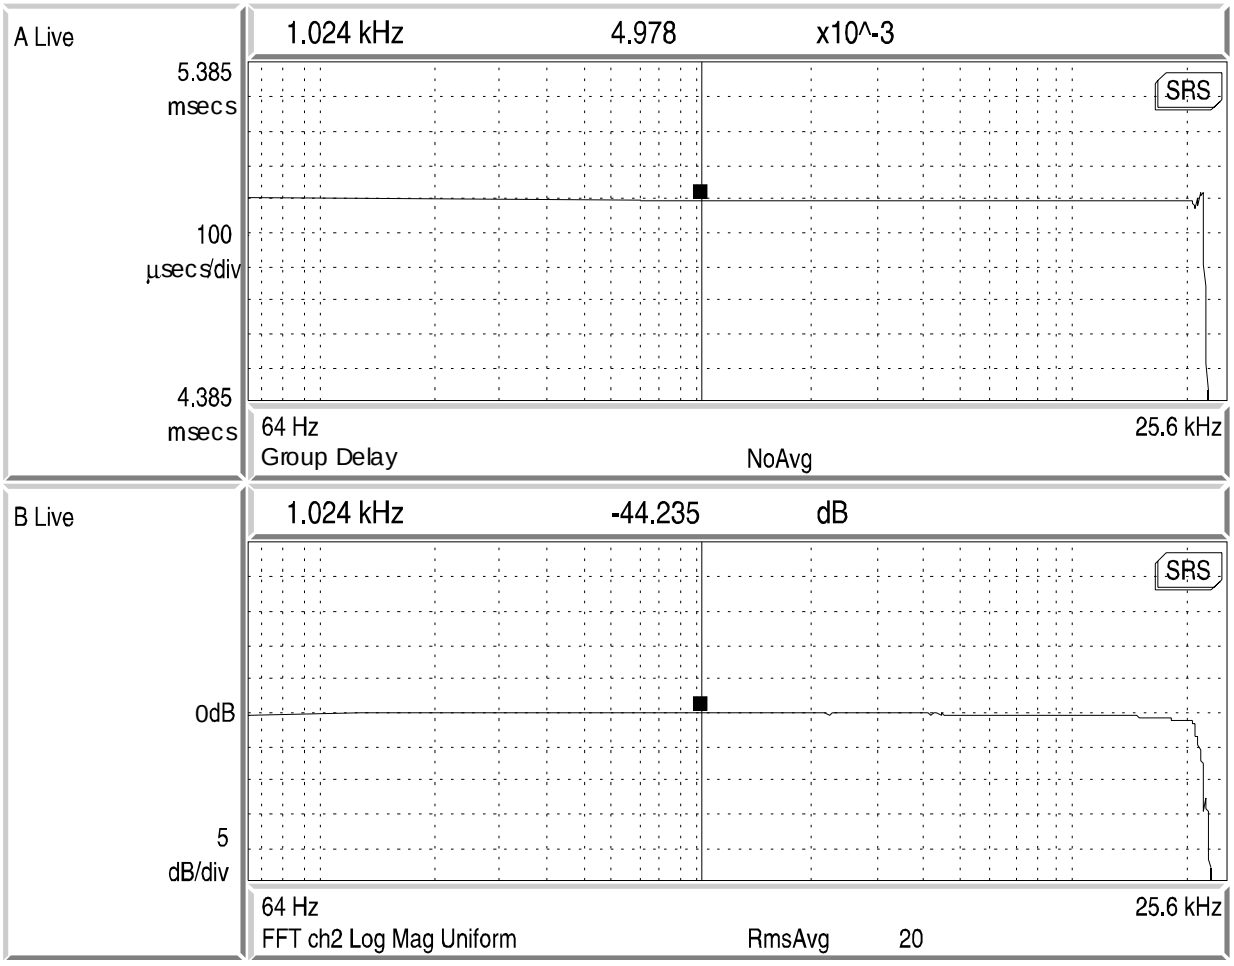

Figure 5: ***20-Band Digital Audio Equalizer*** **Flat Response.**

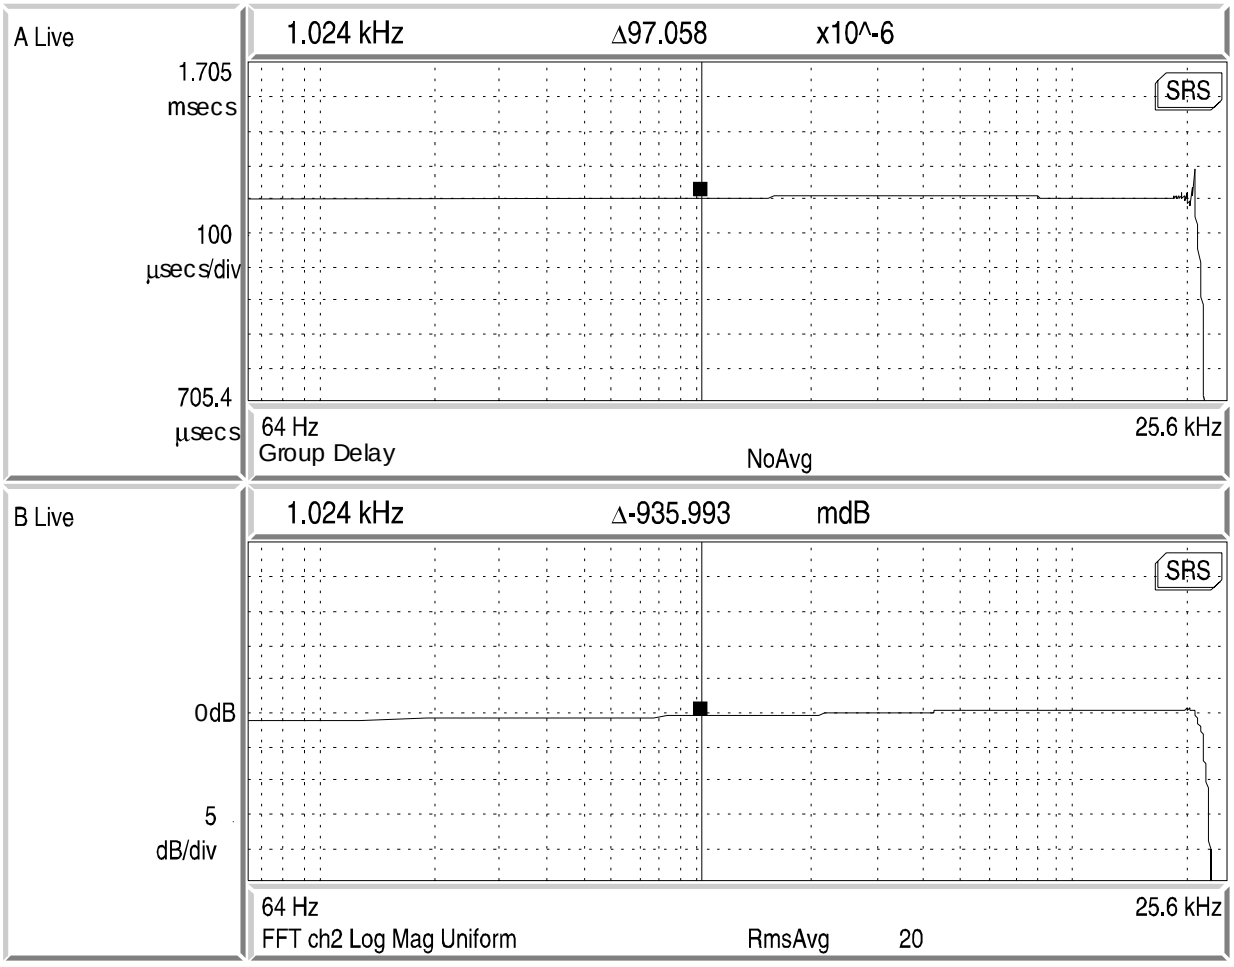

Figure 6: ***Equalizer X*** **Flat Response.**

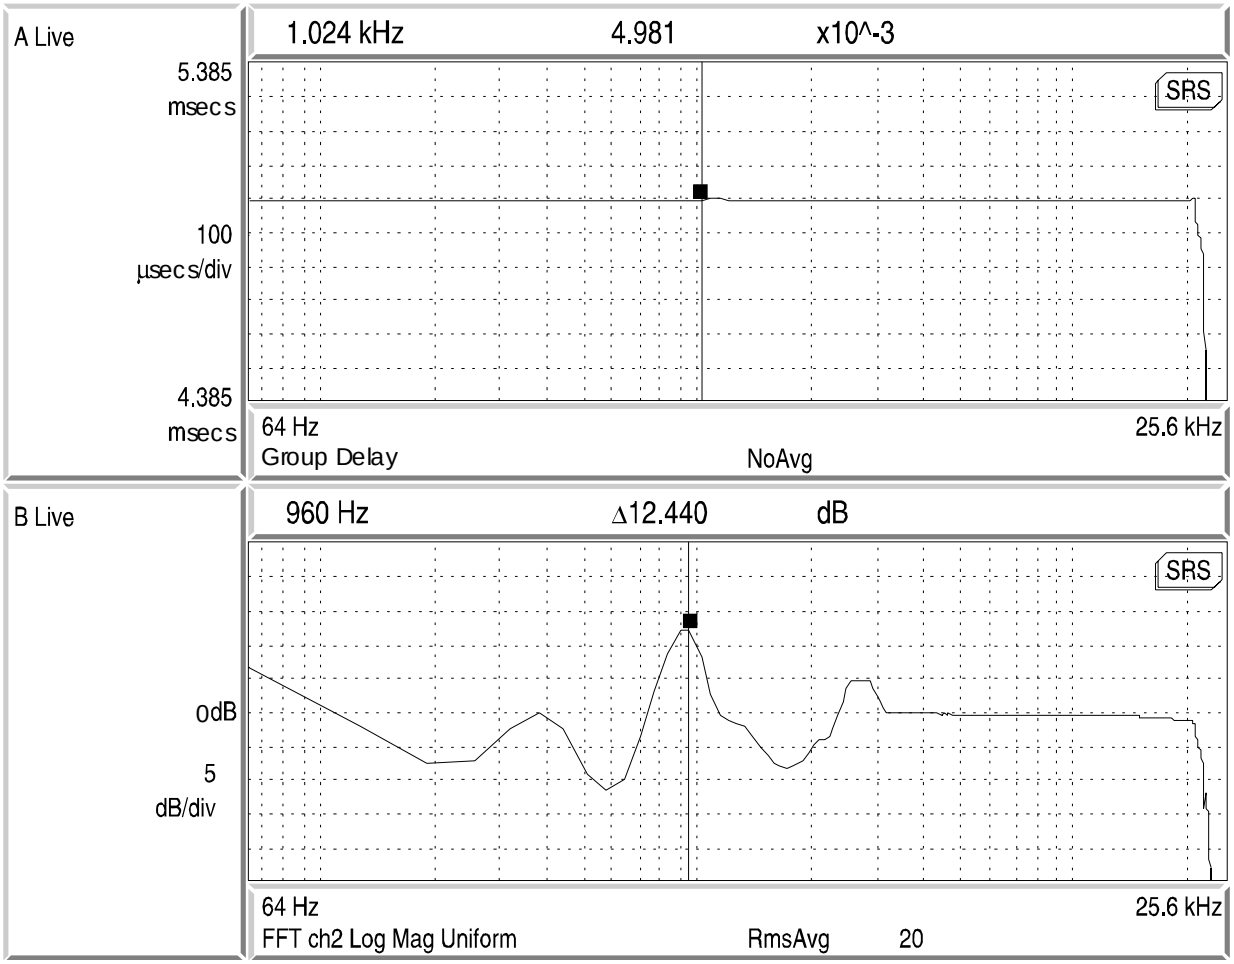

Figure 7: ***20-Band Digital Audio Equalizer*** **Response to Table 2 equalization curve.**

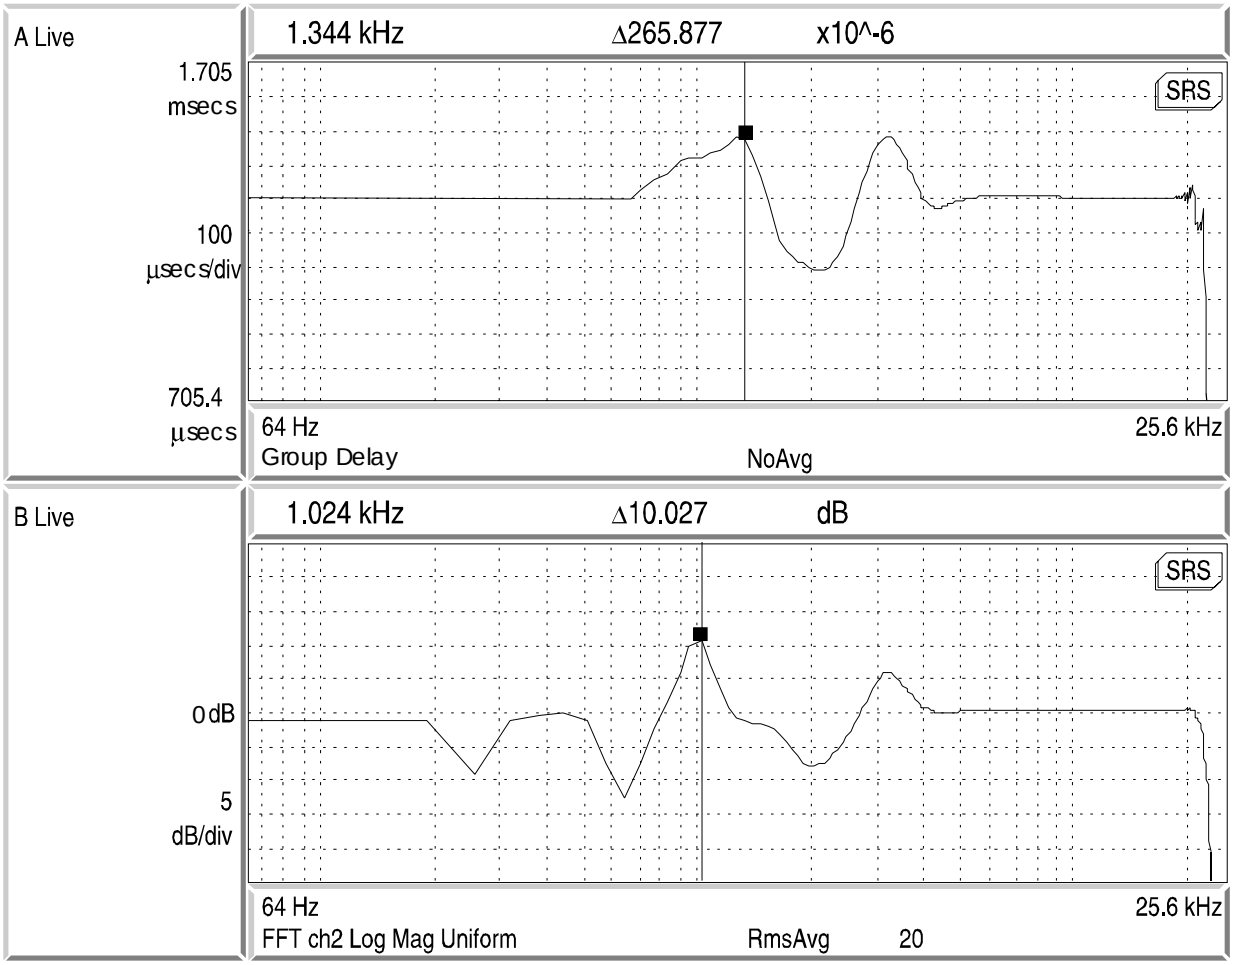

Figure 8: ***Equalizer X*** **Response to Table 2 equalization curve.**